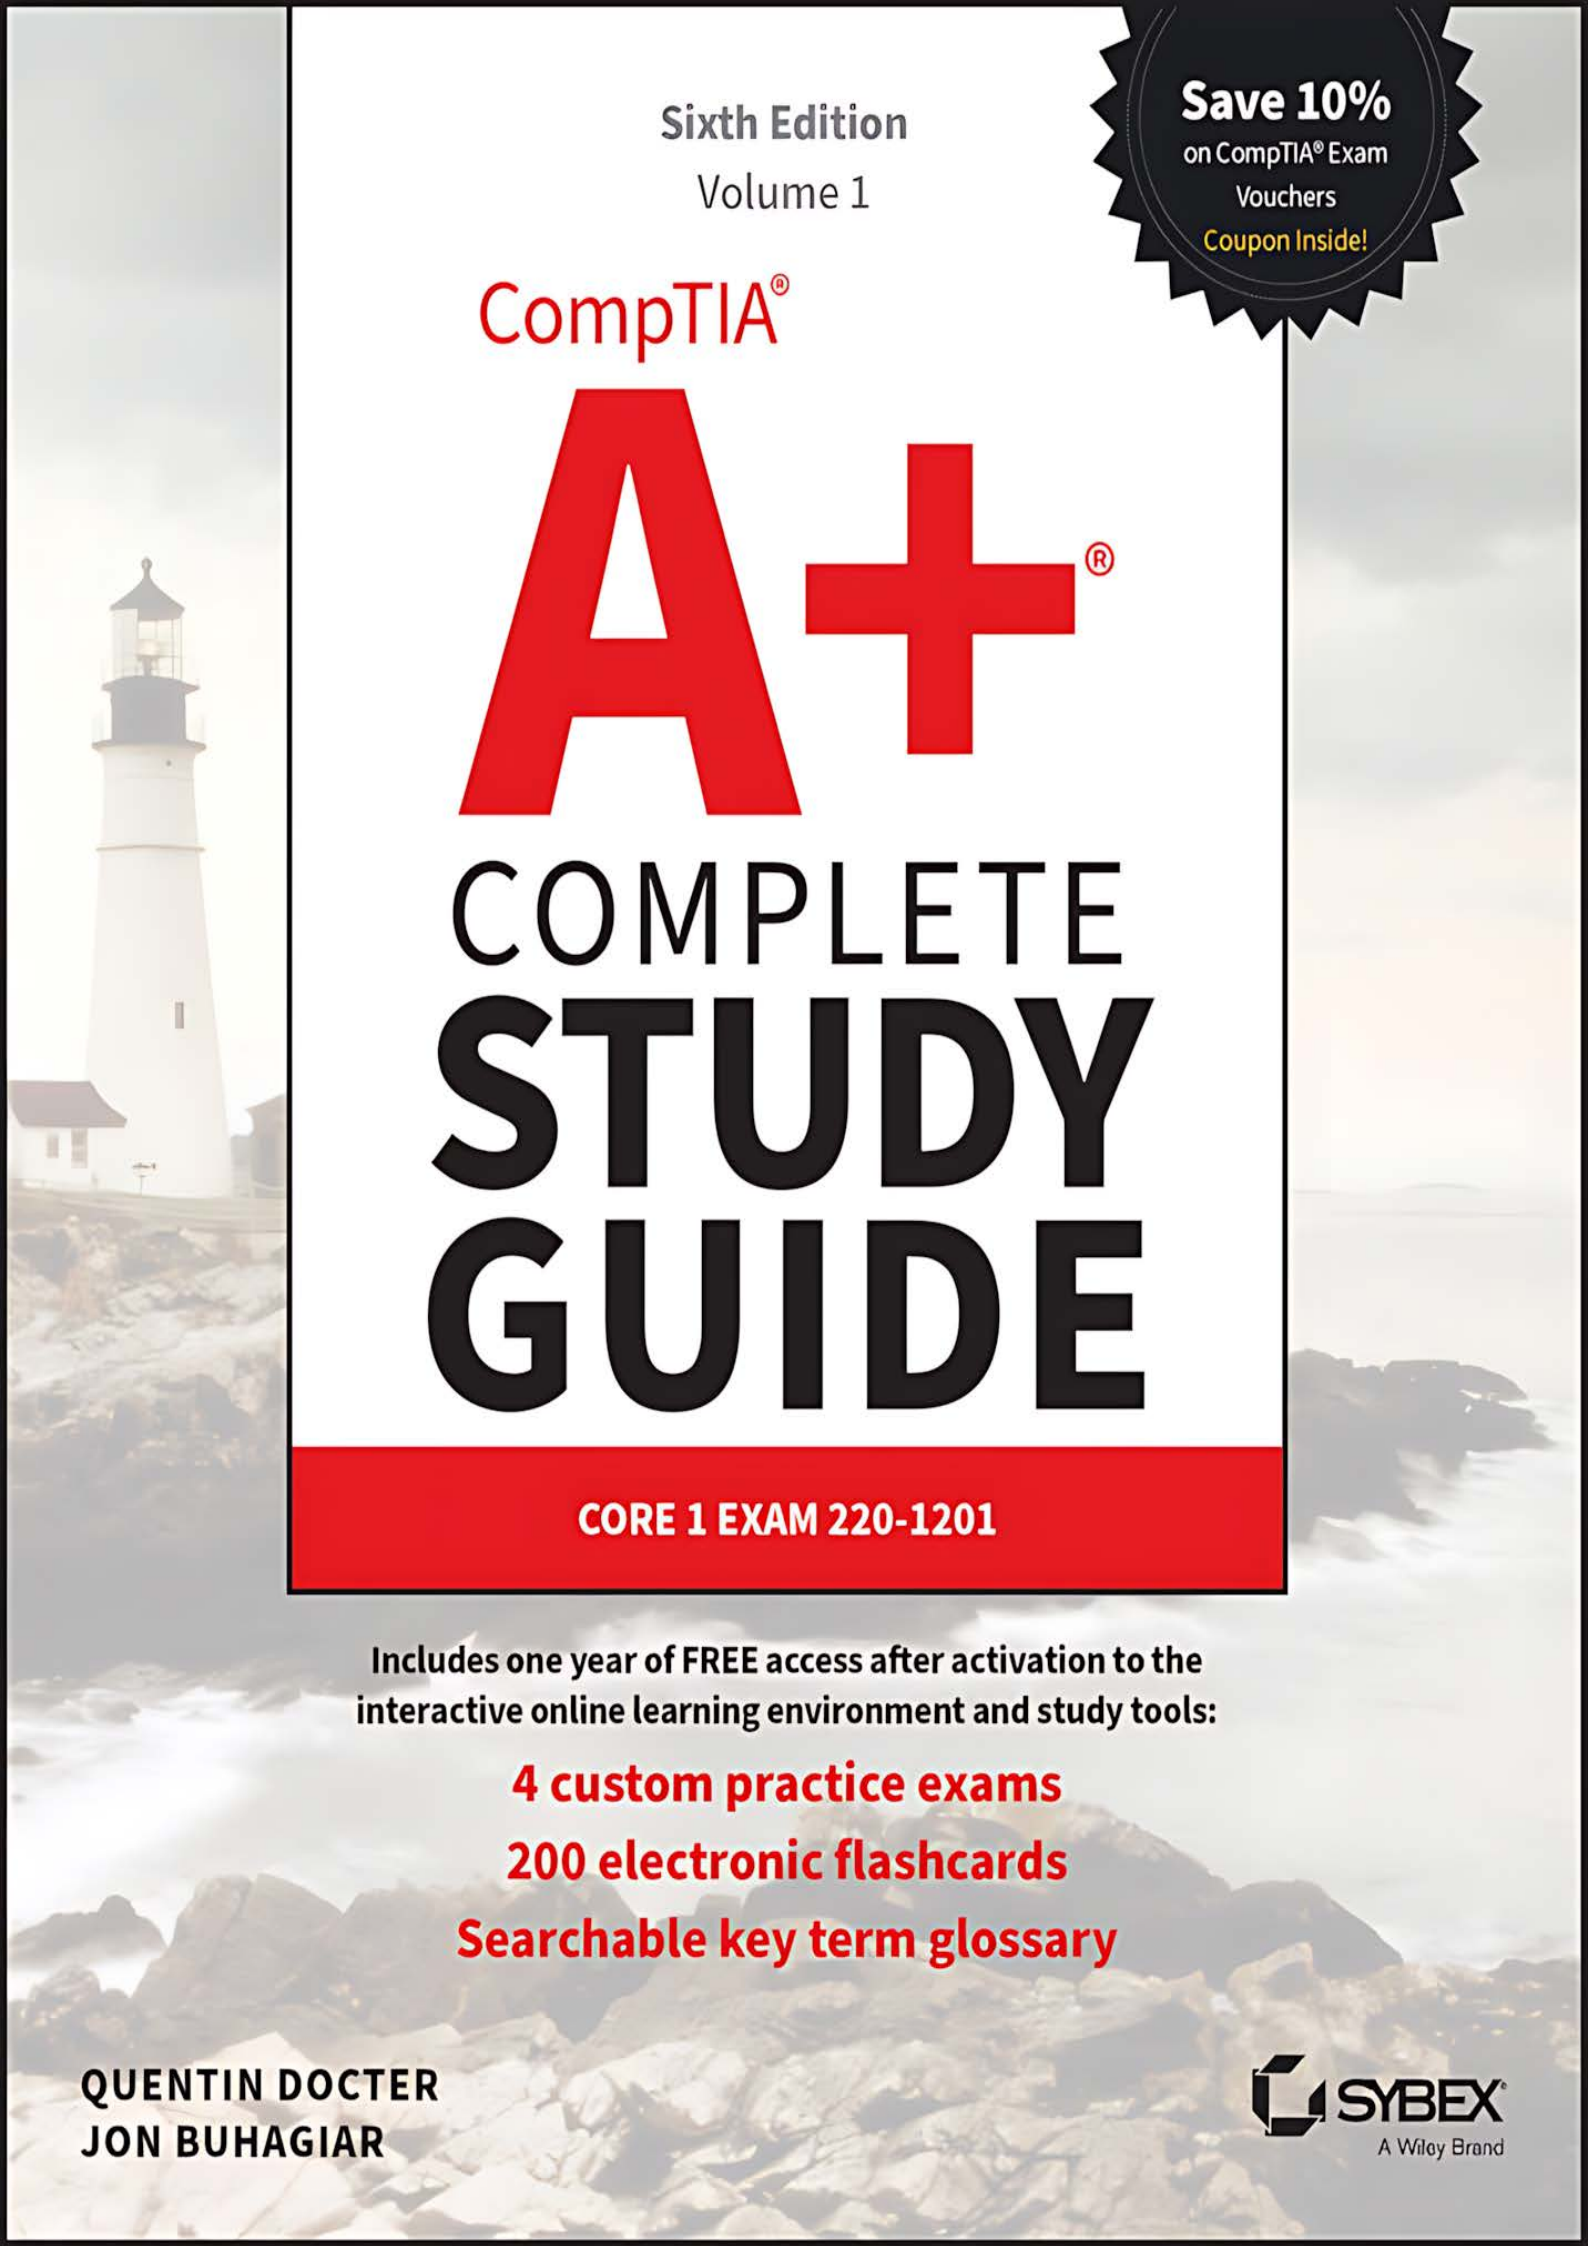
Save 10%
on CompTIA® Exam
Vouchers
Coupon Inside!
Sixth Edition
Volume 1
CompTIA®
A+®
COMPLETE
STUDY
GUIDE
CORE 1 EXAM 220-1201
Includes one year of FREE access after activation to the
interactive online learning environment and study tools:
4 custom practice exams
200 electronic flashcards
Searchable key term glossary
QUENTIN DOCTER
JON BUHAGIAR
SYBEX
A Wiley Brand

Take the Next Step in Your IT Career

Save 10%

on Exam Vouchers*

(up to a $35 value)

*Some restrictions apply. See web page for details.

CompTIA®

Use coupon code WILEY10 during checkout. Redeeming the coupon code is easy:

1. Go to www.comptiastore.com.
2. Browse Certification Vouchers and select the exam voucher you want.
3. Add the voucher to the cart (note that for A+ you will need a separate voucher for each exam).
4. Enter the code WILEY10 on the purchase screen, click Apply and then click Proceed to Checkout to continue and complete the payment process.

SYBEX®

CompTIA®

A+® Complete

Study Guide

Sixth Edition

Volume 1: Core 1 Exam 220-1201

CompTIA®

A+® Complete

Study Guide

Volume 1: Core 1 Exam 220-1201

Sixth Edition

Quentin Docter

Jon Buhagiar

SYBEX®
A Wiley Brand

Copyright © 2025 by John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.
Published simultaneously in Canada and the United Kingdom.

ISBNs for the 2-volume set: 9781394330034 (paperback), 9781394330058 (ePDF), 9781394330041 (ePub)

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at www.wiley.com/go/permission.

The manufacturer's authorized representative according to the EU General Product Safety Regulation is Wiley-VCH GmbH, Boschstr. 12, 69469 Weinheim, Germany, e-mail: Product_Safety@wiley.com.

Trademarks: Wiley, the Wiley logo, and Sybex are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries and may not be used without written permission. CompTIA, A+, and the CompTIA logo are trademarks or registered trademarks of CompTIA, Inc. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.

Limit of Liability/Disclaimer of Warranty: While the publisher and authors have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Further, readers should be aware that websites listed in this work may have changed or disappeared between when this work was written and when it is read. Neither the publisher nor authors shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002. For product technical support, you can find answers to frequently asked questions or reach us via live chat at https://sybexsupport.wiley.com.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Control Number: 2025938071

Cover Design: Wiley
Cover Image: © Jeremy Woodhouse/Getty Images

For the fam.
—Quentin Docter

For my wife and son.
—Jon Buhagiar

Acknowledgments

As we were putting together this book, I was reminded of the proverb that begins "It takes a village . . ." That beginning definitely holds true for creating a book of this scope and size. From beginning to end, scores of dedicated professionals have focused on delivering the best book possible to you, the readers.

First, I need to thank my coauthor, Jon Buhagiar. I appreciate him diving in and dedicating himself to helping to produce an excellent book. I also need to give special thanks to our technical editor, Chris Crayton. He was meticulous and thorough, and challenged me to always find new and better ways to communicate complex concepts. His entire focus was on providing the best training material possible, and I doubt there's better in the business. Now, on to the rest of the team.

Kenyon Brown and Kim Wimpsett kept us on track and moving forward, which was a challenge at times. Saravanan Dakshinamurthy had the fun job of keeping us organized, which is akin to herding cats. Copyeditor Liz Welch reminded me yet again that I am no master of the English language and saved me from butchering it (too badly). Many thanks also go out to our proofreader, Nancy Carrasco, and our indexer, Tom Dinse. Without their great contributions, this book would not have made it into your hands.

On a personal note, I need to thank my family. My wife and kids are all incredibly supportive. Unfortunately, book writing as a side hustle while holding down a full-time job takes up a lot of time. I end up hiding in my office a lot, but they're always there for me, and I couldn't do it without them. Another huge thanks goes to my late grandpa, Joe, who got me into computers and taught me so many lessons I will never be able to repay. Finally, thanks to my friends who keep me relatively sane—Sean, Kurtis, Tim, John, Cory, Kevin, and others—and laugh at me when I tell them I spent my weekend writing about network cables and protocols.

—Quentin Docter

I would like to first thank my coauthor, Quentin Docter. Throughout the writing of this book, he helped me with his insight and his expertise in writing technical books. Without his words of wisdom and guidance, this book would not be the product it stands to be. I would also like to give special thanks to our technical editor, Chris Crayton. His thorough review of the material helped to identify many areas for us to elaborate on and polish.

I would also like to thank the many people who made this book possible: Kenyon Brown at Wiley Publishing, for giving me the opportunity to write this book and work with this wonderful team; Kim Wimpsett, for keeping us on track during the writing process; Pete Gaughan, who also kept us on track and organized during the publishing process; and our copyeditor, Liz Welch, for helping me use proper English. I'd also like to thank the many other people I've never met but who worked behind the scenes to make this book a success.

During the writing of this book, many others in my life supported me, and I'd like to take this opportunity to thank them as well. First and foremost, thanks to my wife and son for their support during the many evenings and weekends spent in my office at the computer. Finally, thanks to my friend and coworker, Bill, for encouraging me daily with his insight and jokes, as well as to all my other coworkers and friends. Thank you.

—Jon Buhagiar

About the Authors

Quentin Docter (A+, Network+, Tech+, Cloud Essentials +, MCSE, CCNA) is an IT author and consultant who started as a tech support agent for Packard Bell in 1994. Since then, he's worked in tech support, network administration, consulting, training, web development, and project management. During his career, he has achieved certifications from CompTIA, Microsoft, Cisco, Novell, and Sun Microsystems. He has written more than a dozen books for Sybex, including books on A+, Tech+, Cloud Essentials+, Server+, Windows, and Solaris 9 certifications, as well as PC hardware and maintenance.

Jon Buhagiar (Network+, A+, CCNA, MCSA, MCSE, BS/ITM) is an IT professional with two decades of experience in higher education. During the past 22 years he has been responsible for network operations at Pittsburgh Technical College and has led several projects, such as virtualization (server and desktop), VoIP, Microsoft 365, and many other projects supporting the quality of education at the college. He has achieved several certifications from Cisco, CompTIA, and Microsoft, and he has taught many of the certification paths. He is the author of several books, including Sybex's *CompTIA Network+ Review Guide: Exam N10-008* (2021) and *CCNA Certification Practice Tests: Exam 200-301* (2020).

About the Technical Editor

Chris Crayton is a technical consultant, trainer, author, and industry-leading technical editor. He has worked as a computer technology and networking instructor, information security director, network administrator, network engineer, and PC specialist. Chris has authored several print and online books on PC repair, CompTIA A+, CompTIA Security+, and Microsoft Windows. He has also served as technical editor and content contributor on numerous technical titles for several of the leading publishing companies. He holds numerous industry certifications, has been recognized with many professional and teaching awards, and has served as a state-level SkillsUSA final competition judge.

Contents at a Glance

Introduction		*xxvii*
Assessment Test for Exam 220-1201		*lv*
Answers to Assessment Test 220-1201		*lxiv*
Volume 1	**CompTIA A+ Core 1 Exam (220-1201)**	**1**
Chapter 1	Motherboards, Processors, and Memory	3
Chapter 2	Expansion Cards, Storage Devices, and Power Supplies	79
Chapter 3	Peripherals, Cables, and Connectors	141
Chapter 4	Laptop and Mobile Device Hardware	191
Chapter 5	Printers and Multifunction Devices	255
Chapter 6	Networking Fundamentals	323
Chapter 7	Introduction to TCP/IP	389
Chapter 8	Network Services	437
Chapter 9	Wireless and SOHO Networking	487
Chapter 10	Mobile Connectivity and Application Support	561
Chapter 11	Virtualization and Cloud Computing	647
Chapter 12	Core Hardware and Network Troubleshooting	697
Chapter 13	Mobile Device, Display, and Printer Troubleshooting	771
Appendix A	Answers to Review Questions	833
Appendix B	Answers to Performance-Based Questions	859
Index		*869*

Contents

Introduction *xxvii*

Assessment Test for Exam 220-1201 *lv*

Answers to Assessment Test 220-1201 *lxiv*

Volume 1 CompTIA A+ Core 1 Exam (220-1201) 1

Chapter 1 Motherboards, Processors, and Memory 3

Understanding Motherboards 6
- Motherboard Form Factors 6
- System Board Components 8

Understanding Processors 42
- CPU Architecture 43
- CPU Characteristics 46

Understanding Memory 50
- Important Memory Terms 50
- Types of Memory 53
- Memory Packaging 59

Understanding Cooling Systems 61
- Case Cooling and Fans 63
- CPU Cooling 66
- Cooling Other Components 70

Summary 71
Exam Essentials 71
Review Questions 73
Performance-Based Question 1 77
Performance-Based Question 2 77

Chapter 2 Expansion Cards, Storage Devices, and Power Supplies 79

Installing and Configuring Expansion Cards 81
- Video 82
- Multimedia 84
- Network Interface Card 86
- Input/Output 88
- Adapter Configuration 89

Understanding Storage Devices 94
- Conventional Hard Disk Drives 94
- Solid-State Drives 99
- RAID 109

Removable Storage and Media 111
Installing, Removing, and Configuring Storage Devices 116
Understanding Power Supplies 119
Power Supply Input 120
Power Supply Output and Wattage Ratings 121
Energy Efficiency 122
Power Connectors 123
Modular Power Supplies 126
Redundant Power Supplies 127
Replacing Power Supplies 130
AC Adapters as Power Supplies 131
Summary 132
Exam Essentials 132
Review Questions 134
Performance-Based Question 139

Chapter 3 Peripherals, Cables, and Connectors 141

Understanding Peripheral Devices 143
Video Devices 144
Audio Devices 157
Input and Output Devices 159
Storage Devices 162
Understanding Cables and Connectors 162
Peripheral Cables and Connectors 163
Serial Ports 172
Video Cables and Connectors 173
Hard Drive Cables and Connectors 181
Summary 184
Exam Essentials 185
Review Questions 186
Performance-Based Question 190

Chapter 4 Laptop and Mobile Device Hardware 191

Working with Laptop and Mobile Device Hardware 194
Understanding the Differences Between Device Types 194
Disassembling and Reassembling Laptops 199
Installing and Configuring Laptop Hardware 202
Setting Up and Configuring Accessories and Ports 244
Summary 249
Exam Essentials 249
Review Questions 250
Performance-Based Question 254

Chapter 5 Printers and Multifunction Devices 255

Understanding Print Technologies and Imaging Processes 257
Impact Printers 258
Inkjet Printers 260
Laser Printers 269
Thermal Printers 281
Installing and Maintaining Printers 282
Installing and Sharing Local Printers 282
Installing and Sharing Networked Printers 298
Performing Appropriate Printer Maintenance 310
Summary 315
Exam Essentials 316
Review Questions 318
Performance-Based Question 322

Chapter 6 Networking Fundamentals 323

Understanding Networking Principles 325
Network Types 326
Primary Network Components 333
Network Resource Access 337
Network Topologies 340
Rules of Communication 344
Identifying Common Network Hardware 348
Network Interface Cards 348
Cables and Connectors 353
Networking Components 371
Summary 381
Exam Essentials 382
Review Questions 384
Performance-Based Question 388

Chapter 7 Introduction to TCP/IP 389

Understanding TCP/IP 391
TCP/IP Structure 392
Understanding IP Addressing 403
IPv6 420
Understanding Virtual Networks 423
Virtual Local Area Networks 423
Virtual Private Networks 425
Summary 429
Exam Essentials 430
Review Questions 432
Performance-Based Question 436

Chapter 8 Network Services 437

Understanding Network Services 439
Server Roles 439
Internet Appliances 465
Legacy/Embedded Systems 473
Internet of Things Devices 475
Summary 480
Exam Essentials 480
Review Questions 482
Performance-Based Question 486

Chapter 9 Wireless and SOHO Networking 487

Understanding Wireless Networking Technologies 489
Understanding 802.11 Networking 490
Bluetooth Networking 505
Radio Frequency Networking Standards 508
Installing and Configuring SOHO Networks 511
Keys to Planning a Network 511
Choosing an Internet Connection 514
Choosing Internal Network Connections 526
Installing the Network Infrastructure 530
Configuring Wireless Routers and Access Points 539
Summary 554
Exam Essentials 554
Review Questions 556
Performance-Based Question 560

Chapter 10 Mobile Connectivity and Application Support 561

Understanding Mobile Connectivity 563
Understanding Cellular Networking Standards 565
Using Cellular Data Connections 567
Establishing Wi-Fi Connectivity 577
Establishing Bluetooth Connectivity 593
Understanding Mobile App Support 606
Understanding Location Services 606
Understanding Mobile Device and App Management 613
Configuring Mobile Device Synchronization 626
Summary 639
Exam Essentials 639
Review Questions 641
Performance-Based Question 645

Chapter 11 Virtualization and Cloud Computing 647

Understanding Virtualization 649
The Purpose of Virtual Machines 650

The Hypervisor 652
Containers 655
Client-Side Virtualization Requirements 656
Understanding Cloud Computing 668
Types of Clouds 668
Cloud Services 669
Important Cloud Characteristics 671
Using Cloud-Based Services 674
Summary 688
Exam Essentials 689
Review Questions 691
Performance-Based Question 695

Chapter 12 Core Hardware and Network Troubleshooting 697

Understanding Troubleshooting Best Practices 701
Troubleshooting Tricks and Tips 701
Structured Troubleshooting Methodology 702
Troubleshooting Motherboards, RAM, CPUs**, and Power Problems** 704
Identifying General Hardware Symptoms and Causes 705
Identifying BIOS/UEFI and POST Routine Problems 710
Identifying Motherboard and CPU Problems 714
Identifying Memory Issues 716
Identifying Power Supply Problems 719
Troubleshooting Drives and RAID **Arrays** 722
Lights and Sounds 723
Devices Not Found 724
Performance Issues 725
S.M.A.R.T. Diagnostics 730
RAID Issues 733
Optical Drive Issues 734
Troubleshooting Networking Problems 735
Using Network Troubleshooting Tools 735
Resolving Connectivity Issues 756
Summary 763
Exam Essentials 763
Review Questions 766
Performance-Based Question 770

Chapter 13 Mobile Device, Display, and Printer Troubleshooting 771

Troubleshooting Video, Projector, and Display Issues 774
Video Input Issues 775
Video Image Problems 776

Other Display Issues 781
Troubleshooting Common Mobile Device Issues 783
Power and Heat Issues 784
Input Problems 786
Connectivity Issues 788
Physical Damage and Malware 792
Other Mobile Device Issues 796
Troubleshooting Printer Issues 799
Impact Printer Issues 800
Inkjet Printer Problems 803
Laser Printer Problems 807
Managing Print Jobs 817
Summary 823
Exam Essentials 824
Review Questions 827
Performance-Based Question 831

Appendix A Answers to Review Questions 833

Chapter 1: Motherboards, Processors, and Memory 834
Chapter 2: Expansion Cards, Storage Devices, and Power Supplies 836
Chapter 3: Peripherals, Cables, and Connectors 838
Chapter 4: Laptop and Mobile Device Hardware 839
Chapter 5: Printers and Multifunction Devices 841
Chapter 6: Networking Fundamentals 843
Chapter 7: Introduction toTCP/IP 845
Chapter 8: Network Services 847
Chapter 9: Wireless and SOHO Networking 849
Chapter 10: Mobile Connectivity and Application Support 851
Chapter 11: Virtualization and Cloud Computing 853
Chapter 12: Core Hardware and Network Troubleshooting 854
Chapter 13: Mobile Device, Display, and Printer Troubleshooting 856

Appendix B Answers to Performance-Based Questions 859

Chapter 1: Motherboards, Processors, and Memory 860
Chapter 2: Expansion Cards, Storage Devices, and Power Supplies 861
Chapter 3: Peripherals, Cables, and Connectors 861
Chapter 4: Laptop and Mobile Device Hardware 862
Chapter 5: Printers and Multifunction Devices 863
Chapter 6: Networking Fundamentals 863
Chapter 7: Introduction to TCP/IP 864

Chapter 8: Network Services 865
Chapter 9: Wireless and SOHO Networking 866
Chapter 10: Mobile Connectivity and Application Support 866
Chapter 11: Virtualization and Cloud Computing 867
Chapter 12: Core Hardware and Network Troubleshooting 867
Chapter 13: Mobile Device, Display, and Printer Troubleshooting 868

Index *869*

Table of Exercises

Exercise 1.1 Entering the UEFI from Windows 11 34
Exercise 2.1 Updating a Device Driver in Windows 11. 91
Exercise 2.2 Removing an Internal Storage Device 117
Exercise 2.3 Installing an Internal Storage Device. 118
Exercise 2.4 Removing a Power Supply 131
Exercise 3.1 Changing the Refresh Rate in Windows 11 147
Exercise 3.2 Changing the Settings for Multiple Monitors 151
Exercise 4.1 Removing Speakers from a Laptop 204
Exercise 4.2 Removing the Display Assembly 210
Exercise 4.3 Removing the Display Panel. 211
Exercise 4.4 Removing the Motherboard from a Laptop. 214
Exercise 4.5 Replacing Laptop Memory 218
Exercise 4.6 Removing an M.2 SSD from a Laptop 222
Exercise 4.7 Removing a Laptop Keyboard 225
Exercise 4.8 Disabling a Touchpad in Windows 11 227
Exercise 4.9 Disabling a Track Point in Windows 11 229
Exercise 4.10 Removing an Internal Laptop Battery 234
Exercise 4.11 Removing the System Fan 236
Exercise 4.12 Removing the CPU Heat Sink 237
Exercise 4.13 Removing the Wireless NIC 238
Exercise 4.14 Removing the CMOS Battery 240
Exercise 4.15 Flashing the System BIOS 241
Exercise 5.1 Identifying the Parts of an Inkjet Printer 267
Exercise 5.2 Installing a USB Printer in Windows 11 297
Exercise 5.3 Installing a TCP/IP Printer in Windows 11 299
Exercise 5.4 Determining if Bonjour Is Installed in Windows 301
Exercise 5.5 Scanning a Document to Google Drive 308
Exercise 5.6 Using an Inkjet Cleaning Solution 314
Exercise 6.1 Pricing Network Cables. 371
Exercise 7.1 Changing IP Address Assignment in Windows 11. 403
Exercise 7.2 Setting Up Windows 11 VPN 427
Exercise 8.1 Configuring Windows 11 to Use a Proxy Server 470
Exercise 9.1 The Cost of Networking 529

Exercise 9.2 Installing an Internal NIC in Windows 11 530
Exercise 10.1 Connecting an iPhone to a Wi-Fi Network 578
Exercise 10.2 Connecting an Android Phone to a Wi-Fi Network 580
Exercise 10.3 Disabling Cellular Use for Data Networking on an iPhone 582
Exercise 10.4 Disabling Cellular Use for Data Networking in Android OS 585
Exercise 10.5 Setting Up a VPN in Android 590
Exercise 10.6 Pairing an Android Phone with a Windows Laptop 594
Exercise 10.7 Configuring Location Services in iOS 610
Exercise 10.8 Configuring Location Services in Android 611
Exercise 10.9 Email Account Configuration on an iPhone 616
Exercise 10.10 Email Account Configuration in Android 619
Exercise 11.1 Enabling Hyper-V in Windows 11 653
Exercise 11.2 Installing VirtualBox and Lubuntu on Windows 11 661
Exercise 11.3 Using Google's Cloud Services 676
Exercise 11.4 Creating a Free AWS Account 678
Exercise 11.5 Creating a Virtual Server in AWS 681
Exercise 12.1 Troubleshooting Practice 721
Exercise 12.2 Viewing and Managing Storage Space 726
Exercise 12.3 Renewing an IP Address in Windows 11 743
Exercise 12.4 Renewing an IP Address from the Command Line 745
Exercise 12.5 Using the `net share` Command in Windows 751
Exercise 13.1 Stopping and Restarting the Print Spooler in Windows 11 820

Introduction

Welcome to the *CompTIA® A+® Complete Study Guide Volume 1: Core 1 Exam 220-1201 Sixth Edition*. This is the sixth edition of our best-selling study guide for the A+ certification sponsored by CompTIA (Computing Technology Industry Association). Thank you for choosing us to help you on your journey toward certification!

This book is written at an intermediate technical level; we assume that you already know how to *use* a personal computer and its basic peripherals, such as USB devices and printers, but we also recognize that you may be learning how to *service* some of that computer equipment for the first time. The exams cover basic computer service topics as well as more advanced issues, and they cover topics that anyone already working as a technician should be familiar with. The exams are designed to test you on these topics in order to certify that you have enough knowledge to fix and upgrade some of the most widely used types of personal computers and operating systems.

In addition to the information in the chapters, we've included a lot of extra material to help your study prep. At the end of each chapter is a list of exam essentials to know as well as 20 review questions to give you a taste of what it's like to take the exams. In addition, there are eight bonus exams of at least 75 questions each. Finally, there are flashcards designed to help your recall. Before you dive into those, though, we recommend you take the assessment test at the end of this introduction to gauge your current knowledge level.

Don't just study the questions and answers—the questions on the actual exams will be different from the practice ones included with this book. The exams are designed to test your knowledge of a concept or objective, so use this book to learn the objective behind the question. That said, we're confident that if you can do well on our quizzes, you will be well equipped to take the real exam.

This book covers more than just the exams, however. We believe in providing our students with a foundation of IT knowledge that will prepare them for real jobs, not just to pass a test. After all, life is not a multiple-choice test with the answers clearly laid out in front of you!

For experienced IT professionals, you can use the book to fill in the gaps in your current computer service knowledge. You may find, as many PC technicians have, that being well versed in all the technical aspects of hardware and operating systems is not enough to provide a satisfactory level of support—you must also have customer-relations skills, understand safety concepts, and be familiar with change management and environmental impacts and controls. We include helpful hints in all of these areas.

What Is CompTIA A+ Certification?

The A+ certification program was developed by CompTIA to provide an industrywide means of certifying the competency of computer service technicians. The A+ certification is granted to those who have attained the level of knowledge and troubleshooting skills that are needed

to provide capable support in the field of personal computers. It is similar to other certifications in the computer industry, such as the Cisco Certified Support Technician (CCST) program and the Microsoft certification programs. The theory behind these certifications is that if you need to have service performed on any of their products, you would sooner call a technician who has been certified in one of the appropriate certification programs than just call the first "expert" in the phone book.

The A+ certification program was created to offer a wide-ranging certification, in the sense that it is intended to certify competence with personal computers and mobile devices from many different makers/vendors. You must pass two tests to become A+ certified:

- The Core 1 220-1201 exam, which covers basic computer concepts, PC hardware, basic networking, mobile devices, and hardware troubleshooting
- The Core 2 220-1202 exam, which covers operating systems, security, software troubleshooting, and operational procedures

You don't have to take the 220-1201 and the 220-1202 exams at the same time. However, the A+ certification is not awarded until you've passed both tests.

Why Become A+ Certified?

There are several good reasons to get your A+ certification. The CompTIA Candidate's Information packet lists five major benefits:

- It demonstrates proof of professional achievement.
- It increases your marketability.
- It provides greater opportunity for advancement in your field.
- It is increasingly a requirement for some types of advanced training.
- It raises customer confidence in you and your company's services.

Provides Proof of Professional Achievement The A+ certification is a status symbol in the computer service industry. Organizations that include members of the computer service industry recognize the benefits of A+ certification and push for their members to become certified. And more people every day are putting the "A+ Certified Technician" emblem on their business cards.

Increases Your Marketability A+ certification makes individuals more marketable to potential employers. A+ certified employees also may receive a higher base salary because employers won't have to spend as much money on vendor-specific training.

Provides Opportunity for Advancement Most raises and advancements are based on performance. A+ certified employees work faster and more efficiently and are thus more productive. The more productive employees are, the more money they make for their company. And, of course, the more money they make for the company, the more valuable they are to the company. So, if an employee is A+ certified, their chances of being promoted are greater.

Fulfills Training Requirements Most major computer hardware vendors recognize A+ certification. Some of these vendors apply A+ certification toward prerequisites in their own respective certification programs, which has the side benefit of reducing training costs for employers.

Raises Customer Confidence As the A+ Certified Technician moniker becomes better known among computer owners, more of them will realize that the A+ technician is more qualified to work on their computer equipment than a noncertified technician.

How to Become A+ Certified

A+ certification is available to anyone who passes two exams. You don't have to work for any particular company. It's not a secret society. It is, however, an elite group. To become A+ certified, you must do two things:

- Pass the Core 1 220-1201 exam.
- Pass the Core 2 220-1202 exam.

The exams can be taken at any Pearson VUE testing center. If you pass both exams, you will get a certificate in the mail from CompTIA saying that you have passed, and you will also receive a lapel pin and business card.

To register for the tests, go to `www.pearsonvue.com/comptia`. You'll be asked for your name, Social Security number (an optional number may be assigned if you don't wish to provide your Social Security number), mailing address, phone number, employer, when and where you want to take the test, and your credit card number. (Payment arrangements must be made at the time of registration.)

Tips for Taking the A+ Exam

Here are some general tips for taking your exam successfully:

- Bring two forms of ID with you. One must be a photo ID, such as a driver's license. The other can be a major credit card or a passport. Both forms must include a signature.
- Arrive early at the exam center so that you can relax and review your study materials, particularly tables and lists of exam-related information. When you enter the testing room, you will need to leave everything outside; you won't be able to bring any materials into the testing area.
- Read the questions carefully. Don't be tempted to jump to an early conclusion. Make sure that you know exactly what each question is asking.
- Don't leave any unanswered questions. Unanswered questions are scored against you. Some questions will have multiple correct responses. When a question has more than one correct answer, a message at the bottom of the screen will prompt you to, for example, choose two. Be sure to read the messages displayed to know how many correct answers you must choose.

- When answering multiple-choice questions that you're not sure about, use a process of elimination to get rid of the obviously incorrect answers first. Doing so will improve your odds if you need to make an educated guess.
- For the hands-on questions where you need to complete a task, double-check that you have finished all parts of it, including checking all tabs available to you if they give you an interface to click on. Also, understand that these questions will take longer than multiple-choice ones. Spend the time needed on them; you will be able to make up time on the shorter multiple-choice questions.
- On form-based tests (nonadaptive), because the hard questions will take the most time, save them for last. You can move forward and backward through the exam.
- For the latest pricing on the exams and updates to the registration procedures, visit CompTIA's website at `www.comptia.org`.

Who Should Read This Book?

If you are one of the many people who want to pass the A+ exams, and pass them confidently, then you should buy this book and use it to study for the exams.

This book was written to prepare you for the challenges of the real IT world, not just to pass the A+ exams. This study guide will do that by describing in detail the concepts on which you'll be tested.

What Does This Book Cover?

CompTIA A+ Complete Study Guide: Exam 220-1201 and Exam 220-1202 covers everything you need to know to pass the CompTIA A+ exams.

Volume 1 has 13 chapters, and it covers all the topics on which you will be tested for Exam 220-1201:

Chapter 1: Motherboards, Processors, and Memory Chapter 1 details the characteristics of motherboards and their built-in components. The CPU, RAM, and BIOS, which are attached to the motherboard, are also presented in Chapter 1.

Chapter 2: Expansion Cards, Storage Devices, and Power Supplies Chapter 2 presents internal expansion options and popular forms of storage devices in use today, including traditional hard drives, solid-state drives, flash drives, and memory cards. Capacities, form factors, and the makeup of these components are also discussed. Finally, this chapter looks at providing power to computer components.

Chapter 3: Peripherals, Cables, and Connectors Chapter 3 covers installation and characteristics of external peripheral devices, the ports they use, cables required, and the connectors used to attach them.

Chapter 4: Laptop and Mobile Device Hardware Chapter 4 covers topics such as mobile device–specific hardware, components within a laptop display, and mobile device features.

Chapter 5: Printers and Multifunction Devices Chapter 5 starts by discussing different printing technologies, such as impact, inkjet, laser, and thermal printers. It then moves on to cover installing printers and performing printer maintenance.

Chapter 6: Networking Fundamentals Chapter 6 covers characteristics of networking cable types and connectors, network devices, networking concepts, and network components.

Chapter 7: Introduction to TCP/IP Chapter 7 details the most common network protocol in use today. It covers TCP/IP structure, addressing (including IPv6), common protocols in the suite, and virtual networks.

Chapter 8: Network Services Chapter 8 focuses on the various services used in networks. It covers the roles of DNS, DHCP, Syslog, file servers, and print servers. In addition, it covers Internet appliances, legacy/embedded systems, and Internet of Things (IoT) devices.

Chapter 9: Wireless and SOHO Networking Chapter 9 contains two main sections. The first is on wireless networking standards and security, and the second discusses setting up a small office, home office (SOHO) network and choosing an Internet connection type.

Chapter 10: Mobile Connectivity and Application Support Chapter 10 details the specific features of mobile devices that are unique to their formats. This chapter provides extensive hands-on steps for configuring a variety of connectivity options and application support on these devices.

Chapter 11: Virtualization and Cloud Computing Chapter 11 starts off by covering virtualization concepts, as well as the components related to virtualization such as hypervisors and containers. The chapter then covers the concepts and components of cloud computing.

Chapter 12: Core Hardware and Network Troubleshooting Chapter 12 introduces some tips and tricks for resolving computer problems, which every technician should know. It then covers troubleshooting core hardware issues related to the motherboard, CPU, RAM, and power. Then, it explores troubleshooting storage drives and RAID arrays, as well as networking tools and troubleshooting.

Chapter 13: Mobile Device, Display, and Printer Troubleshooting Chapter 13 finishes the troubleshooting section, including video issues, mobile device issues, and printers.

Volume 2 includes 10 chapters, and it covers all the topics on which you will be tested for Exam 220-1202.

What's Included in the Book

We've included several learning tools throughout the book:

Assessment Tests At the end of this introduction is an assessment test that you can use to check your readiness for the exam. Take this before you start reading the book; it will help you determine the areas on which you might need to brush up. The answers to the assessment test questions appear on a separate page after the last question of the test. Each answer includes an explanation and the chapter the material appears in.

Objective Map and Opening List of Objectives In this introduction, we have included a detailed exam objective map that shows you where each exam objective is covered. In addition, each chapter opens with a list of the exam objectives it covers. Use these resources to see exactly where each exam topic is covered.

Exam Essentials Each chapter, just after the summary, includes a number of "exam essentials." These are the key topics that you should focus on when preparing for the exam.

Chapter Review Questions To test your knowledge as you progress through the book, each chapter ends with review questions. As you finish each chapter, answer the review questions and then check your answers—the correct answers and explanations are in Appendix A. You can go back to reread the section that deals with each question you got wrong to ensure that you answer correctly the next time that you're tested on the material.

Performance-Based Questions The A+ exams may ask you to perform tasks to correctly answer a question. This may mean dragging and dropping items into the correct order or clicking the correct area of an image or screenshot. While it's hard to simulate those in a paper book, the end of each chapter features a performance-based question. These questions will have you think through a topic presented in the chapter, and then perform a task. The correct answers are in Appendix B.

Interactive Online Learning Environment and Test Bank

The interactive online learning environment that accompanies *CompTIA A+ Complete Study Guide: Exam 220-1201 and Exam 220-1202* provides a test bank with study tools to help you prepare for the certification exams and increase your chances of passing them the first time! The test bank includes the following elements:

Sample Tests All the questions and answers in this book are provided, including the assessment tests, which you'll find at the end of this introduction, and the review questions at the end of each chapter. In addition, there are eight practice exams. Use these questions to test your knowledge of the study guide material. The online test bank runs on multiple devices.

Flashcards Four sets of questions are provided in digital flashcard format (a question followed by a single correct answer). You can use the flashcards to reinforce your learning and provide last-minute test prep before the exam.

Glossary The key terms from this book and their definitions are available as a fully searchable PDF.

eBooks Read eBook versions of this Study Guide in multiple formats on your favorite device.

Audio Instruction Over 2 hours of recorded instruction from the authors to help you study for the certification exam with confidence.

Go to `www.wiley.com/go/sybextestprep` to register and gain access to this interactive online learning environment and test bank with study tools.

Like all exams, the CompTIA A+ certification is updated periodically and may eventually be retired or replaced. At some point after CompTIA is no longer offering this exam, the old editions of our books and online tools will be retired. If you have purchased this book after the exam was retired, or are attempting to register in the Sybex online learning environment after the exam was retired, please know that we make no guarantees that this exam's online Sybex tools will be available once the exam is no longer available.

How to Use This Book

If you want a solid foundation for preparing for the A+ exams, this is the book for you. We've spent countless hours putting together this book with the intention of helping you prepare for the exams.

This book is loaded with valuable information, and you will get the most out of your study time if you understand how we put the book together. Here's a list that describes how to approach studying:

1. Take the assessment tests immediately following this introduction. It's okay if you don't know any of the answers—that's what this book is for. Carefully read over the explanations for any questions you get wrong, and make note of the chapters where that material is covered.
2. Study each chapter carefully, making sure you fully understand the information and the exam objectives listed at the beginning of the chapter. Again, pay extra-close attention to any chapter that includes material covered in questions you missed on the assessment test.

3. Read over the Summary and Exam Essentials. These will highlight the sections from the chapter you need to be familiar with before sitting for the exam.
4. Answer all the review questions at the end of each chapter. Specifically note any questions that confuse you, and study the corresponding sections of the book again. Don't just skim these questions. Make sure that you understand each answer completely.
5. Go over the electronic flashcards. These help you to prepare for the latest A+ exams, and they're really great study tools.
6. Take the practice exams.

Performance-Based Questions

CompTIA includes performance-based questions on the A+ exams. These are not the traditional multiple-choice questions with which you're probably familiar. These questions require the candidate to know how to perform a specific task or series of tasks. The candidate will be presented with a scenario and will be asked to complete a task. They will be taken to a simulated environment where they will have to perform a series of steps and will be graded on how well they complete the task.

The Sybex test engine does not include performance-based questions. However, at the end of each chapter, we have included a section called "Performance-Based Question," which is designed to measure how well you understood the chapter's topics. Some simply ask you to complete a task for which there is only one correct response. Others are more subjective, with multiple ways to complete the task. We provide the most logical or practical solution in Appendix B. Note that these questions may cover topic areas not covered in the actual A+ performance-based questions. However, we feel that being able to think logically is a great way to learn.

The CompTIA A+ Exam Objectives for Core 1 220-1201

The A+ exams consist of the Core 1 220-1201 exam and the Core 2 220-1202 exam. The following are the detailed exam objectives for the Core 1 220-1201 exam.

Exam objectives are subject to change at any time without prior notice and at CompTIA's sole discretion. Please visit the A+ Certification page of CompTIA's website (`https://comptia.org/certifications/a`) for the most current listing of exam objectives.

A+ Certification Exam Objectives: 220-1201

The following table lists the domains measured by this examination and the extent to which they are represented on the exam:

Domain	Percentage of exam
1.0 Mobile Devices	13%
2.0 Networking	23%
3.0 Hardware	25%
4.0 Virtualization and Cloud Computing	11%
5.0 Hardware and Network Troubleshooting	28%
Total	100%

Objective Map

The following table lists where you can find the objectives covered in this book.

Objective	Chapter(s)
1.0 Mobile Devices	
1.1 Given a scenario, monitor mobile device hardware and use appropriate replacement techniques.	4

- Battery
- Keyboard/keys
- Random-access memory (RAM)
- Hard disk drive (HDD)/solid-state drive (SSD)
- Wireless cards
- Physical privacy and security components
  - Biometrics
  - Near-field scanner feature
- Wi-Fi antenna connector/placement
- Camera/webcam
- Microphone

Objective	Chapter(s)
1.2 Compare and contrast accessories and connectivity options for mobile devices.	4

▪ Connection methods
  ▪ Universal Serial Bus (USB)/USB-C/microUSB/miniUSB
  ▪ Lightning
  ▪ Near-field communication (NFC)
  ▪ Bluetooth
  ▪ Tethering/hotspot
▪ Accessories
  ▪ Stylus
  ▪ Headsets
  ▪ Speakers
  ▪ Webcam
▪ Docking station
▪ Port replicator
▪ Trackpad/drawing pad/track points	

1.3 Given a scenario, configure basic mobile device network connectivity and provide application support.	10

▪ Wireless/cellular data network (enable/disable)
  ▪ 3G/4G/5G
  ▪ Hotspot
  ▪ Wi-Fi
  ▪ Subscriber Identity Module (SIM)/eSIM	

Objective	Chapter(s)

▪ Bluetooth
  ▪ Enable Bluetooth
  ▪ Enable pairing
  ▪ Find a device for pairing
  ▪ Enter the appropriate personal identification number (PIN) code
  ▪ Test connectivity
▪ Location services
  ▪ Global positioning system (GPS) services
  ▪ Cellular location services
▪ Mobile device management (MDM)
  ▪ Device configurations
  ▪ Corporate
  ▪ Bring your own device (BYOD)
  ▪ Policy enforcement
  ▪ Corporate applications
▪ Mobile device synchronization
  ▪ Recognizing data caps
  ▪ Calendar
  ▪ Contacts
  ▪ Business applications
  ▪ Mail
  ▪ Cloud storage	

Objective	Chapter(s)
2.0 Networking	
2.1 Compare and contrast Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) ports, protocols, and their purposes.	7

▪ Ports and protocols
  ▪ 20-21 – File Transfer Protocol (FTP)
  ▪ 22 – Secure Shell (SSH)
  ▪ 23 – Telnet
  ▪ 25 – Simple Mail Transfer Protocol (SMTP)
  ▪ 53 – Domain Name System (DNS)
  ▪ 67/68 – Dynamic Host Configuration Protocol (DHCP)
  ▪ 80 – Hypertext Transfer Protocol (HTTP)
  ▪ 110 – Post Office Protocol 3 (POP3)
  ▪ 143 – Internet Mail Access Protocol (IMAP)
  ▪ 137-139 Network Basic Input/Output System (NetBIOS)/NetBIOS over TCP/IP (NetBT)
  ▪ 389 – Lightweight Directory Access Protocol (LDAP)
  ▪ 443 – Hypertext Transfer Protocol Secure (HTTPS)
  ▪ 445 – Server Message Block (SMB)/Common Internet File System (CIFS)
  ▪ 3389 – Remote Desktop Protocol (RDP)
▪ TCP vs. UDP	

2.2 Explain wireless networking technologies.	9

▪ Frequencies
  ▪ 2.4GHz
  ▪ 5GHz
  ▪ 6GHz
▪ Channels
  ▪ Regulations
  ▪ Channel selection	

Objective	Chapter(s)
▪ Widths ▪ Frequencies ▪ Bands ▪ Bluetooth ▪ 802.11 standards ▪ NFC ▪ Radio-frequency identification (RFID)	
2.3 Summarize services provided by networked hosts. ▪ Server roles ▪ DNS ▪ DHCP ▪ Fileshare ▪ Print servers ▪ Mail servers ▪ Syslog ▪ Web servers ▪ Authentication, Authorization, and Accounting (AAA) ▪ Database servers ▪ Network Time Protocol (NTP) ▪ Internet appliances ▪ Spam gateways ▪ Unified threat management (UTM) ▪ Load balancers ▪ Proxy servers ▪ Legacy/embedded systems ▪ Supervisory control and data acquisition (SCADA) ▪ Internet of Things (IoT) devices	8

Objective	Chapter(s)
2.4 Explain common network configuration concepts.	7, 8

▪ DNS
  ▪ A
  ▪ AAAA
  ▪ Canonical Name (CNAME)
  ▪ Mail exchanger (MX)
  ▪ Text (TXT)
  ▪ Spam management
  ▪ DomainKeys Identified Mail (DKIM)
  ▪ Sender Policy Framework (SPF)
  ▪ Domain-based Message Authentication, Reporting, and Conformance (DMARC)	

▪ DHCP
  ▪ Leases
  ▪ Reservations
  ▪ Scope
  ▪ Exclusions	

▪ Virtual LAN [local area network] (VLAN)	
▪ Virtual private network (VPN)	
2.5 Compare and contrast common networking hardware devices.	6
▪ Routers	

▪ Switches
  ▪ Managed
  ▪ Unmanaged	

Objective	Chapter(s)
▪ Access points	
▪ Patch panel	
▪ Firewall	
▪ Power over Ethernet (PoE)	
▪ Injectors	
▪ Switch	
▪ PoE standards	
▪ Cable modem	
▪ Digital subscriber line (DSL)	
▪ Optical network terminal (ONT)	
▪ Network interface card (NIC)	
▪ Physical media access control (MAC) address	
2.6 Given a scenario, configure basic wired/wireless small office/home office (SOHO) networks.	7
▪ Internet Protocol (IP) addressing	
▪ IPv4	
▪ Private addresses	
▪ Public addresses	
▪ IPv6	
▪ Automatic Private IP Addressing (APIPA)	
▪ Static	
▪ Dynamic	
▪ Subnet mask	
▪ Gateway	

Objective	Chapter(s)
2.7 Compare and contrast internet connection types, network types, and their characteristics.	6, 9

▪ Internet connection types
  ▪ Satellite
  ▪ Fiber
  ▪ Cable
  ▪ DSL
  ▪ Cellular
  ▪ Wireless internet service provider (WISP)
▪ Network types
  ▪ LAN
  ▪ Wide area network (WAN)
  ▪ Personal area network (PAN)
  ▪ Metropolitan area network (MAN)
  ▪ Storage area network (SAN)
  ▪ Wireless local area network (WLAN)	

2.8 Explain networking tools and their purposes.	12
▪ Crimper ▪ Cable stripper ▪ Wi-Fi analyzer ▪ Toner probe ▪ Punchdown tool ▪ Cable tester ▪ Loopback plug ▪ Network tap	

Objective	Chapter(s)
3.0 Hardware	
3.1 Compare and contrast display components and attributes.	3

- Types
  - Liquid crystal display (LCD)
    - In-plane switching (IPS)
    - Twisted nematic (TN)
    - Vertical alignment (VA)
  - Organic light-emitting diode (OLED)
  - Mini light-emitting diode (Mini-LED)
- Touch screen/digitizer
- Inverter
- Attributes
  - Pixel density
  - Refresh rates
  - Screen resolution
  - Color gamut

3.2 Summarize basic cable types and their connectors, features, and purposes. 3, 6

- Network cables
  - Copper
    - Categories
      - T568A/T568B standards
    - Coaxial
    - Shielded twisted pair
      - Direct burial
    - Unshielded twisted pair
  - Plenum-rated
  - Optical
    - Single-mode
    - Multimode

Objective	Chapter(s)

- Peripheral cables
  - USB 2.0
  - USB 3.0
  - Serial
  - Thunderbolt
- Video cables
  - High-definition Multimedia Interface (HDMI)
  - DisplayPort
  - Digital Visual Interface (DVI)
  - Video Graphics Array (VGA)
  - USB-C
- Hard drive cables
  - Serial Advanced Technology Attachment (SATA)
  - External SATA (eSATA)
- Adapters
- Connector types
  - RJ11
  - RJ45
  - F-type
  - Straight tip (ST)
  - Subscriber connector (SC)
  - Lucent connector (LC)
  - Punchdown block
  - MicroUSB
  - MiniUSB
  - USB-C
  - Molex
  - Lightning
  - DB9

Objective	Chapter(s)
3.3 Compare and contrast RAM characteristics.	1

▪ Form factors
  ▪ Small Outline Dual In-line Memory Module (SODIMM)
  ▪ Dual In-line Memory Module (DIMM)
▪ Double Data Rate (DDR) iterations
▪ Error-correcting code (ECC) vs. non-ECC RAM
▪ Channel configurations	

3.4 Compare and contrast storage devices.	2

▪ Hard drives
  ▪ Spindle speeds
  ▪ Form factors
  ▪ 2.5-inch
  ▪ 3.5-inch
▪ Solid-state drives
  ▪ Communications interfaces
  ▪ Non-volatile Memory Express (NVMe)
  ▪ SATA
  ▪ Peripheral Component Interconnect Express (PCIe)
  ▪ Serial Attached SCSI [Small Computer System Interface] (SAS)
  ▪ Form factors
  ▪ M.2
  ▪ Mini-serial Advanced Technology Attachment (mSATA)
▪ Drive configurations
  ▪ Redundant Array of Independent Disks (RAID) 0, 1, 5, 6, 10
▪ Removable storage
  ▪ Flash drives
  ▪ Memory cards
▪ Optical drives	

Objective	Chapter(s)
3.5 Given a scenario, install and configure motherboards, central processing units (CPUs), and add-on cards.	1, 2

- Motherboard form factors
  - Advanced Technology Extended (ATX)
  - microATX
  - Information Technology eXtended (ITX)
- Motherboard connector types
  - Peripheral Component Interconnect (PCI)
  - PCIe
  - Power connectors
  - SATA
  - eSATA
  - Headers
  - M.2
- Motherboard compatibility
  - CPU socket types
    - Advanced Micro Devices, Inc. (AMD)
    - Intel
  - Multisocket
- BIOS/Unified Extensible Firmware Interface (UEFI) settings
  - Boot options
  - USB permissions
  - Trusted Platform Module (TPM) security features
  - Fan considerations

Objective	Chapter(s)

- Secure Boot
- Boot password
- BIOS password
- Temperature monitoring

- Virtualization support
- Encryption
  - TPM
  - Hardware security module (HSM)
- CPU architecture
  - x86/x64
  - Advanced RISC [Reduced Instruction Set Computer] Machine (ARM)
  - Core configurations
- Expansion cards
  - Sound card
  - Video card
  - Capture card
  - Network interface card
- Cooling
  - Fans
  - Heat sink
  - Thermal paste/pads
  - Liquid

Objective	Chapter(s)
3.6 Given a scenario, install the appropriate power supply.	2
▪ Input 110–120 VAC vs. 220–240 VAC ▪ Output 3.3V vs. 5V vs. 12V ▪ 20+4 pin motherboard connector ▪ Redundant power supply ▪ Modular power supply ▪ Wattage rating ▪ Energy efficiency	
3.7 Given a scenario, deploy and configure multifunction devices/ printers and settings.	5

▪ Properly unbox device and consider set-up location
▪ Use appropriate drivers for a given operating system
  ▪ Printer Control Language (PCL) vs. postscript
▪ Firmware
▪ Device connectivity
  ▪ USB
  ▪ Ethernet
  ▪ Wireless
▪ Public/shared devices
  ▪ Printer share
  ▪ Print server
▪ Configuration settings
  ▪ Duplex
  ▪ Orientation
  ▪ Tray settings
  ▪ Quality
▪ Security
  ▪ User authentication	

Objective	Chapter(s)
▪ Badging ▪ Audit logs ▪ Secured prints ▪ Network scan services ▪ Email ▪ SMB ▪ Cloud services ▪ Automatic document feeder (ADF)/flatbed scanner	
3.8 Given a scenario, perform appropriate printer maintenance.	5
▪ Laser ▪ Maintenance: Replace toner, apply maintenance kit, calibrate, and clean ▪ Inkjet ▪ Ink cartridge, printhead, roller, and feeder ▪ Maintenance: Clean printheads, replace cartridges, calibrate, and clear jams ▪ Thermal ▪ Feed assembly ▪ Special thermal paper ▪ Maintenance: Replace paper, clean heating element, and remove debris ▪ Impact ▪ Multipart paper ▪ Maintenance: Replace ribbon, printhead, and paper	

Objective	Chapter(s)
4.0 Virtualization and Cloud Computing	
4.1 Explain virtualization concepts.	11

▪ Purpose of virtual machines
  ▪ Sandbox
  ▪ Test development
  ▪ Application virtualization
  ▪ Legacy software/OS
  ▪ Cross-platform virtualization
▪ Requirements
  ▪ Security
  ▪ Network
  ▪ Storage
▪ Desktop virtualization
  ▪ Virtual Desktop Infrastructure (VDI)
▪ Containers
▪ Hypervisors
  ▪ Type 1
  ▪ Type 2	

4.2 Summarize cloud computing concepts.	11

▪ Common cloud models
  ▪ Private cloud
  ▪ Public cloud
  ▪ Hybrid cloud
  ▪ Community cloud
  ▪ Infrastructure as a service (IaaS)	

Objective	Chapter(s)
▪ Software as a service (SaaS) ▪ Platform as a service (PaaS) ▪ Cloud characteristics ▪ Shared resources vs. dedicated resources ▪ Metered utilization ▪ Ingress/egress ▪ Elasticity ▪ Availability ▪ File synchronization ▪ Multitenancy	
5.0 Hardware and Network Troubleshooting	
5.1 Given a scenario, troubleshoot motherboards, RAM, CPUs, and power. ▪ Common symptoms ▪ Power-on self-test (POST) beeps ▪ Proprietary crash screens ▪ Blank screen ▪ No power ▪ Sluggish performance ▪ Overheating ▪ Burning smell ▪ Random shutdown ▪ Application crashes ▪ Unusual noise ▪ Capacitor swelling ▪ Inaccurate system date/time	12

Objective	Chapter(s)
5.2 Given a scenario, troubleshoot drive and RAID issues.	12

- Common symptoms
  - Light-emitting diode (LED) status indicators
  - Grinding noises
  - Clicking sounds
  - Bootable device not found
  - Data loss/corruption
  - RAID failure
  - Self-monitoring and Reporting Technology (S.M.A.R.T.) failure
  - Extended read/write times
  - Low performance input/output operations per second (IOPS)
  - Missing drives in OS
  - Array missing
  - Audible alarms

5.3 Given a scenario, troubleshoot video, projector, and display issues.	13

- Common symptoms
  - Incorrect input source
  - Physical cabling issues
  - Burnt-out bulb
  - Fuzzy image
  - Display burn-in
  - Dead pixels
  - Flashing screen
  - Incorrect color display
  - Audio issues
  - Dim image
  - Intermittent projector shutdown
  - Sizing issues
  - Distorted image

Objective	Chapter(s)
5.4 Given a scenario, troubleshoot common mobile device issues.	13

- Common symptoms
  - Poor battery health
  - Swollen battery
  - Broken screen
  - Improper charging
  - Poor/no connectivity
  - Liquid damage
  - Overheating
  - Digitizer issues
  - Physically damaged ports
  - Malware
  - Cursor drift/touch calibration
  - Unable to install new applications
  - Stylus does not work
  - Degraded performance

5.5 Given a scenario, troubleshoot network issues.	12

- Common symptoms
  - Intermittent wireless connectivity
  - Slow network speeds
  - Limited connectivity
  - Jitter
  - Poor Voice over Internet Protocol (VoIP) quality
  - Port flapping
  - High latency
  - External interference
  - Authentication failures
  - Intermittent internet connectivity

Objective	Chapter(s)
5.6 Given a scenario, troubleshoot printer issues.	13

- Lines down the printed pages
- Garbled print
- Paper jams
- Faded prints
- Paper not feeding
- Multipage misfeed
- Multiple prints pending in queue
- Speckling on printed pages
- Double/echo images on the print
- Grinding noise
- Finishing issues
  - Staple jams
  - Hole punch
- Incorrect page orientation
- Tray not recognized
- Connectivity issues
- Frozen print queue

Exam objectives are subject to change at any time without prior notice at CompTIA's sole discretion. Please visit CompTIA's website (www.comptia.org) for the most current listing of exam objectives.

How to Contact the Publisher

If you believe you have found a mistake in this book, please bring it to our attention. At John Wiley & Sons, we understand how important it is to provide our customers with accurate content, but even with our best efforts an error may occur.

In order to submit your possible errata, please email it to our Customer Service Team at wileysupport@wiley.com with the subject line "Possible Book Errata Submission."

Assessment Test for Exam 220-1201

1. Which of the following is not considered a system component that can be found inside a computer?
   - **A.** CPU
   - **B.** RAM
   - **C.** PCIe graphics adapter
   - **D.** Motherboard
2. Which of the following is a physical memory format installed directly in today's desktop computer systems?
   - **A.** DIMM
   - **B.** HDD
   - **C.** SSD
   - **D.** eMMC
3. Which of the following are components that can commonly be found on a motherboard? (Choose all that apply.)
   - **A.** Headers
   - **B.** Fan connectors
   - **C.** Gyroscope
   - **D.** Scanner
   - **E.** HDD
4. Which of the following statements is true regarding ECC RAM?
   - **A.** ECC RAM has full functionality in a non-ECC motherboard.
   - **B.** ECC RAM will not fit in a non-ECC motherboard.
   - **C.** ECC RAM will fit but not work in a non-ECC motherboard.
   - **D.** ECC RAM will fit and function without error-correction in a non-ECC motherboard.
5. Which of the following hard drives will fit into an M.2 slot with a B-key and an M-key?
   - **A.** B-key
   - **B.** M-key
   - **C.** B+M-key
   - **D.** All of the above

6. You are replacing a power supply with a new one. What is the primary characteristic you should look for to ensure all internal components can be supported?
   - **A.** Amperage rating
   - **B.** 20-pin to 24-pin motherboard adapter
   - **C.** Wattage rating
   - **D.** Voltage rating

7. A user has asked you to install a RAID array. Except in the case of RAID 0, which two things do all types of RAID offer? (Choose two.)
   - **A.** Faster read speeds
   - **B.** Faster write speeds
   - **C.** Redundancy
   - **D.** Fault tolerance
   - **E.** Automatic restores from backup after a drive failure

8. You are installing a new SATA hard drive. Which of the following is true, without using an adapter?
   - **A.** It will connect to a SATA controller only.
   - **B.** It will connect to a SATA controller or SAS controller.
   - **C.** It will connect to a SATA controller or PATA controller.
   - **D.** It will connect to a SATA controller or a PCIe controller.

9. Which of the following types of connectors would you expect to see on the end of a serial cable?
   - **A.** DB-9
   - **B.** RS-232
   - **C.** Molex
   - **D.** VGA

10. Which of the following video connectors has an analog variety and a digital variety?
   - **A.** VGA
   - **B.** HDMI
   - **C.** DisplayPort
   - **D.** DVI

11. Which of the following connectors will you not find on the end of a USB cable? (Choose two.)
   - **A.** Lightning
   - **B.** USB-C
   - **C.** Molex
   - **D.** DB-9
   - **E.** eSATA

12. In an LCD panel, what is the name of the component that provides power to the backlight?

A. Diffuser

B. Inverter

C. Converter

D. Digitizer

13. You are installing a new display for a laptop. Which LCD component is responsible for providing brightness?

A. Backlight

B. Inverter

C. Screen

D. Digitizer

14. Your laptop has 8 GB of installed memory and uses shared video memory. If the video card is using 512 MB, how much is left for the rest of the system?

A. 8 GB

B. 7.5 GB

C. 512 MB

D. Cannot determine

15. On a typical laptop, which of the following components runs through the hinge from the base of the laptop up into the display unit?

A. Webcam connector

B. Speaker wires

C. Microphone connector

D. Wi-Fi antenna

16. Which of the following laptop accessories is most likely to have a full-sized internal hard drive bay?

A. Docking station

B. Port replicator

C. SAN

D. NAS

17. A technician is preparing to perform inkjet printer maintenance. Which of the following should they consider?

A. Replace toner, apply maintenance kit, calibrate, clean.

B. Replace paper, clean heating element, remove debris.

C. Replace ribbon, replace printhead, replace paper.

D. Clean printheads, replace cartridges, calibrate, clear jams.

18. What is the function of the laser in a laser printer?

A. It heats up the toner so that it adheres to the page.

B. It charges the paper so that it will attract toner.

C. It creates an image of the page on the drum.

D. It cleans the drum before a page is printed.

19. What is the component called that stores the material that ends up printed to the page in a laser printer?

A. Toner cartridge

B. Ink cartridge

C. Laser module

D. Laser cartridge

20. When used with network scanning services, what does SMB allow for?

A. Scan to the cloud

B. Scan to email

C. Scan to a network folder

D. Scan to secured print

21. Which network connectivity device stops broadcasts from being sent to computers on a different network segment?

A. Hub

B. Switch

C. Router

D. Firewall

22. Which of the following cable types is made with a plastic coating that does not produce toxic fumes when it burns?

A. Plenum

B. Fiber

C. Coaxial

D. UTP

23. Which of the following are standards used by coaxial cable? (Choose two.)

A. RG-6

B. RG-59

C. RJ11

D. RJ45

24. You have a PoE switch that says it supports PoE Type 3. How many watts can it provide to attached devices?

A. 30 watts

B. 60 watts

C. 100 watts

D. 120 watts

25. On which port does FTP run by default?

A. 21

B. 25

C. 53

D. 80

26. Which of the following protocols can be used by a client to access email on a server?

A. DNS

B. FTP

C. SMTP

D. IMAP

27. Which of the following protocols provide secure network transmissions? (Choose two.)

A. SMTP

B. SNMP

C. SSH

D. HTTPS

28. Which of the following IP addresses is not routable on the Internet?

A. 192.16.33.220

B. 172.18.1.101

C. 33.33.33.33

D. 192.186.33.220

29. One of your network users was recently caught at work browsing pornographic websites. Which of the following servers could be installed to prohibit this activity?

A. Web

B. Security

C. Proxy

D. DNS

30. Which of the following Internet appliances will include services provided by an IDS or IPS?

A. UTM

B. Spam gateway

C. Load balancer

D. Proxy server

31. For security purposes, your network needs to ensure that all devices have synchronized time stamps for login events. Which network service will help provide this?

A. DNS

B. DHCP

C. AAA

D. NTP

32. Which DNS record type is used for IPv4 addresses?

A. AAAA

B. A

C. CNAME

D. MX

33. Which of the following is a company that provides direct access to the Internet for home and business computer users?

A. ASP

B. ISP

C. DNS

D. DNP

34. What is the maximum range for a Bluetooth Class 2 device?

A. 100 meters

B. 10 meters

C. 1 meter

D. 0.5 meters

35. Which of the following 802.11 technologies can communicate in the 2.4 GHz and 5 GHz frequencies? (Choose two.)

A. 802.11ax

B. 802.11ac

C. 802.11n

D. 802.11g

36. Which wireless networking technology has the longest range?

A. Bluetooth

B. 802.11

C. NFC

D. RFID

37. Which of the following is used to remotely manage corporate mobile devices on a network?

A. Location services

B. Synchronization

C. BYOD

D. MDM

38. What term refers to copying data between a mobile device and a computer system to mirror such things as contacts, programs, pictures, and music?

A. Calibration

B. Remote wipe

C. Pairing

D. Synchronization

39. You want to send emails from your smartphone using a secure connection. Which port will enable that?

A. 25

B. 110

C. 143

D. 587

40. A network administrator wants ten users in three different countries to be able to access all of the same data for a project. Which solution is best?

A. MDM

B. Data cap

C. Synchronization

D. Cloud storage

41. When using cloud services, which of the following may incur an additional charge as part of metered utilization?

A. Elasticity

B. Availability

C. Egress

D. Ingress

42. Google Docs is an example of what type of cloud service?

A. SaaS

B. IaaS

C. PaaS

D. GaaS

43. Which type of software is required to run client-side virtualization on your home network?

A. Terminal emulation

B. Process replication

C. Hyperthreading

D. Hypervisor

44. What type of cloud has elements from a public cloud run by a CSP and others owned by the client company on its internal network?

A. Private

B. Public

C. Hybrid

D. Community

45. A desktop computer has recently started exhibiting slow performance, particularly when opening and saving files to the hard drive. What should you try first?

A. Replace the hard drive.

B. Remove old files and applications to free up disk space.

C. Delete and reinstall the OS.

D. Defragment the hard drive.

46. A network cable has not been working properly. Another technician suggests testing the cable. Which tool should be used to do this? (Choose two.)

A. Loopback plug

B. Network tap

C. Cable tester

D. Toner probe

47. You are troubleshooting a network issue. What feature do switches most likely have to help troubleshoot connectivity issues?

A. Port toggles

B. Beep codes

C. Indicator lights

D. Log entries

48. When you turn a computer on, you immediately hear a loud screeching sound coming from it. What components could be making this noise? (Choose two.)

A. SSD

B. HDD

C. CPU

D. PSU

49. Your laser printer has recently starting printing vertical white lines on documents it prints. What is the most likely cause of the problem?

A. The print driver is faulty.

B. The fuser is not heating properly.

C. There is toner on the transfer corona wire.

D. There is a scratch on the EP drum.

50. A laptop's display has several spots where it does not light up. If the screen is set to a solid white background, there are at least a dozen dots that will still be black. What's the solution to this problem?

A. Change the screen resolution to a lower one.

B. Replace the laptop display.

C. Degauss the laptop display.

D. Replace the laptop display inverter.

51. Which of the following should a user do to avoid malware on an Android phone?

A. Use a screen protector.

B. Block the digitizer from receiving external input.

C. Install apps from only the Google Play store.

D. None of the above. Android phones are not susceptible to malware.

Answers to Assessment Test 220-1201

1. C. System components are essential for the basic functionality of a computer system. Many of the landmarks found on the motherboard can be considered system components, even expansion slots to a degree. What you plug into those slots, however, must be considered peripheral to the basic operation of the system. See Volume 1 Chapter 1 for more information.

2. A. Except for DIMMs, all options represent some form of secondary storage, all of which are covered in Volume 1 Chapter 2. See Volume 1 Chapter 1 for more information.

3. A, B. Motherboards commonly have headers for lights and buttons and for powering cooling fans. Gyroscopes are most commonly found in mobile devices. Scanners are external devices. Although there might be one or more types of HDD interfaces built into the motherboard, the HDD itself is not. See Volume 1 Chapter 1 for more information.

4. D. Error-correcting code (ECC) RAM will fit and work in a non-ECC motherboard, but it will not have error-correcting capabilities. See Volume 1 Chapter 1 for more information.

5. A. An M.2 slot with B- and M-keys has posts in the B- and M- positions. Therefore, the drive must have notches in both the B- and M- positions or it will not fit. See Volume 1 Chapter 2 for more information.

6. C. Wattage is the unit of output for power supplies. Make sure that the power supply has enough wattage to support all components. See Volume 1 Chapter 2 for more information.

7. C, D. Except for RAID 0, all implementations of RAID offer a way to recover from the failure of at least one drive, which is an example of fault tolerance, through the implementation of some mechanism that stores redundant information for that purpose. Some RAID types offer faster read and/or write performance. RAID 1, for instance does not guarantee either. RAID also does not automatically restore from backup. See Volume 1 Chapter 2 for more information.

8. B. SATA drives will of course connect to SATA controllers. They will connect to SAS controllers as well, although they will be much slower than SAS hard drives. See Volume 1 Chapter 2 for more information.

9. A. Serial cables use a DB-9 connector. Some people will call them an RS-232 connector, but RS-232 is the transmission standard used by serial cables. Molex are power connectors for PATA hard drives and optical drives. VGA is a video connector. See Volume 1 Chapter 3 for more information.

10. D. DVI comes in both analog (DVI-A) and digital (DVI-D) varieties. VGA is always analog; HDMI and DisplayPort are always digital. See Volume 1 Chapter 3 for more information.

11. C, D. USB cables are multipurpose cables that connect many types of peripherals. Connectors include USB-C, Lightning, and eSATA. Molex connectors provide power to PATA hard drives and optical drives. A DB-9 connector is found on the end of a serial cable. See Volume 1 Chapter 3 for more information.

12. B. A backlight is given power by an inverter, which converts DC power from the power supply to AC power for the backlight. There are no diffusers or converters in an LCD panel. A digitizer takes input from a stylus or fingertip and turns it into digital signals. See Volume 1 Chapter 3 for more information.

13. A. The backlight provides light to the LCD screen. The inverter provides power to the backlight, and the screen displays the picture. Digitizers take physical input (such as touch) and turn it into electronic data. See Volume 1 Chapter 4 for more information.

14. B. If the laptop is using shared video memory, then the system memory is shared with the video card. If the video card is using 512 MB (half a gigabyte), then there is 7.5 GB left for the system. See Volume 1 Chapter 4 for more information.

15. D. In most laptops, the Wi-Fi antenna runs from the base of the unit through the hinge and into the display. The webcam and microphone are generally part of the display unit only. Speakers are most often in the base of the laptop, not the display. See Volume 1 Chapter 4 for more information.

16. A. A docking station might have an internal full-sized hard drive bay. Port replicators only have additional ports, such as HDMI, DisplayPort, and USB. Storage area networks (SAN) and network attached storage (NAS) are not laptop accessories. See Volume 1 Chapter 4 for more information.

17. D. Inkjet printer maintenance includes clean printheads, replace cartridges, calibrate, clear jams. Laser printer maintenance includes replace toner, apply maintenance kit, calibrate, clean. Thermal printer maintenance includes replace paper, clean heating element, remove debris. Impact printer maintenance includes replace ribbon, replace printhead, replace paper. See Volume 1 Chapter 5 for more information.

18. C. The laser creates an image on the photosensitive drum that is then transferred to the paper by the transfer corona. The fuser heats up the toner so that it adheres to the page. The transfer corona charges the page, and the eraser lamp cleans the drum before a page is printed. A rubber blade is also used to remove toner physically from the drum. See Volume 1 Chapter 5 for more information.

19. A. Laser printers use toner, which they melt to the page in the image of the text and graphics being printed. A toner cartridge holds the fine toner dust until it is used in the printing process. See Volume 1 Chapter 5 for more information.

20. C. The Server Message Block (SMB) protocol allows for the saving of scanned documents onto a network folder. See Volume 1 Chapter 5 for more information.

21. C. A router does not pass along broadcasts to computers on other segments. Hubs and switches send broadcasts along because they do not segment traffic at the logical network address level. See Volume 1 Chapter 6 for more information.

22. A. The plenum is an air space within buildings used to circulate breathable air for occupants of a building. Fiber, coaxial, and UTP cables are normally enclosed with a PVC plastic coating, which produces toxic fumes if burned. For air ventilation spaces, plenum grade cable is required. All cable types have plenum-rated versions. See Volume 1 Chapter 6 for more information.

23. A, B. RG-6 and RG-59 are coaxial cable standards. RJ11 and RJ45 are connectors used with twisted pair cabling. See Volume 1 Chapter 6 for more information.

24. B. Power over Ethernet Type 3 can provide 60 W of power. PoE is 15.4 W, PoE+ is 30 W, and PoE++ Type 4 is 100 W. See Volume 1 Chapter 6 for more information.

25. A. FTP operates on ports 20 and 21. Port 25 is used by SMTP, port 53 is DNS, and port 80 is HTTP. See Volume 1 Chapter 7 for more information.

26. D. The IMAP and POP3 protocols can be used to retrieve email from mail servers. DNS is used to resolve host names to IP addresses. FTP is for file transfers, and SMTP is used to send email. See Volume 1 Chapter 7 for more information.

27. C, D. Secure Shell (SSH) and Hypertext Transfer Protocol Secure (HTTPS) are both secure transport protocols. Simple Network Management Protocol (SNMP) and Simple Mail Transfer Protocol (SMTP) are not inherently secure. See Volume 1 Chapter 7 for more information.

28. B. Private IP addresses are not routable on the Internet. The Private IP address ranges are 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16. See Volume 1 Chapter 7 for more information.

29. C. A proxy server can be configured to block access to websites containing potentially objectionable material. Web servers host websites. There is no specific security server, although there are authentication servers which provide security. A DNS server resolves hostnames to IP addresses. See Volume 1 Chapter 8 for more information.

30. A. Unified threat management (UTM) is a security-based Internet appliance that will likely contain the services of an intrusion detection system (IDS) or intrusion protection system (IPS). See Volume 1 Chapter 8 for more information.

31. D. A Network Time Protocol (NTP) server can ensure that all devices on the network are operating on the same time. DNS, DHCP, and AAA do not provide timekeeping services. See Volume 1 Chapter 8 for more information.

32. B. IPv4 addresses are identified with an A DNS record type. AAAA is for IPv6, CNAME is an alias, and MX is a mail server. See Volume 1 Chapter 8 for more information.

33. B. An Internet service provider (ISP) provides direct access to the Internet. ASP is a programming language. DNS resolves hostnames to IP addresses. DNP is not a computer networking acronym. See Volume 1 Chapter 9 for more information.

34. B. A Bluetooth class 2 device (which includes headsets, keyboards, and mice) is 10 meters. Class 1 devices can transmit up to 100 meters, class 3 one meter, and class 4 half a meter. See Volume 1 Chapter 9 for more information.

35. A, C. 802.11ax (Wi-Fi 6) and 802.11n were designed to communicate at 2.4 GHz and 5 GHz. 802.11ac runs at 5 GHz, and 802.11g runs at 2.4 GHz. See Volume 1 Chapter 9 for more information.

36. B. 802.11 has the longest range. Bluetooth is next longest, and then the two shortest are the radio frequency standards NFC and RFID. See Volume 1 Chapter 9 for more information.

37. D. Mobile device management (MDM) controls which devices are on the network and can wipe entire devices. Location services does not manage software. Synchronization can back up apps and data but does not help with updates or deleting apps. BYOD is when a user owns the device that they use on a corporate network. For more information, see Volume 1 Chapter 10.

38. D. Synchronizing a mobile device with a computer system allows you to mirror personal data between the devices, regardless of which one contains the most current data. Calibration refers to matching the device's and user's perceptions of where the user is touching the screen. Remote wipes allow you to remove personal data from a lost or stolen device. Pairing is what must be done in Bluetooth for two Bluetooth devices to connect and communicate. For more information, see Volume 1 Chapter 10.

39. D. Port 587 is used by SMTP (an email sending protocol) secured with TLS. Ports 25 (SMTP), 110 (POP), and 143 (IMAP) are not secure. See Volume 1 Chapter 10 for more information.

40. D. Cloud storage is the best solution to allow geographically dispersed users to access the same data. MDM is mobile device management, which configures the devices allowed to join a network. Synchronization is the act of mirroring all unique changes and additions from one device to the other. See Volume 1 Chapter 10 for more information.

41. C. Metered utilization means client usage is tracked and they are billed for their usage. Many service providers charge extra for downloading data, which is data egress (uploading, or ingress, is generally free). Elasticity and availability are separate cloud characteristics from metered utilization. See Volume 1 Chapter 11 for more information.

42. A. Google Docs is software, so it is an example of software as a service (SaaS). IaaS is infrastructure as a service. PaaS is platform as a service. GaaS is not a test objective and is not currently a cloud service type. See Volume 1 Chapter 11 for more information.

43. D. The hypervisor is the key piece of software needed for virtualization. Terminal emulation, process replication, and Hyper-Threading are not related to client-side virtualization. See Volume 1 Chapter 11 for more information.

44. C. A hybrid cloud is a combination of a public cloud (owned and operated by a CSP) and a private cloud (owned and operated by an individual company). A community cloud is one created by multiple organizations with similar interests. See Volume 1 Chapter 11 for more information.

45. B. Hard drive performance can slow down if the drive gets short on space, particularly under 10 percent free space. Try removing old files and applications to free up disk space first. If that doesn't resolve it, then defragment. If neither of those work, it could indicate a failing hard drive. See Volume 1 Chapter 12 for more information.

46. C. Cable testers are for verifying that a cable works properly. Network taps create copies of network traffic for analysis. A loopback plug is for testing the ability of a network adapter to send and receive. A toner probe is used to trace a cable from one location to another. See Volume 1 Chapter 12 for more information.

47. C. Switches and other connectivity devices have indicator lights to help indicate when things are working normally and when something isn't working right. There are no port toggles and switches don't use beep codes. A switch might make a log entry into an event logger, but this is less likely than it having indicator lights. See Volume 1 Chapter 12 for more information.

48. B, D. A screeching sound can be made only by components that have moving parts. This includes conventional hard drives (HDDs) and power supplies (PSUs) with fans. Solid-state drives (SSDs) and CPUs have no moving parts. See Volume 1 Chapter 12 for more information.

49. C. White streaks on printouts are most likely caused by toner on the transfer corona wire. Vertical black lines are caused by a scratch or a groove in the EP drum. If the fuser was not heating properly, toner would not bond to the paper and you would have smearing. Faulty print drivers will cause garbage to print or there will be no printing at all. See Volume 1 Chapter 13 for more information.

50. B. The symptom being described is called dead pixels. If they will not light up, the only solution is to replace the display. Changing the resolution won't help. Degaussing was done on old CRT monitors. The inverter provides power to the laptop display backlight, but if it were failing, the screen would be dim or flickering. See Volume 1 Chapter 13 for more information.

51. C. Android OS can be subjected to malware. To avoid malware, install apps only from the Google Play store and avoid tapping on suspicious links. See Volume 1 Chapter 13 for more information.

CompTIA A+ Core 1 Exam (220-1201)

VOLUME 1

Chapter 1

Motherboards, Processors, and Memory

THE FOLLOWING COMPTIA A+ 220-1201 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

✓ **3.0 Hardware**

✓ **3.3 Compare and contrast RAM characteristics.**

- Form factors
  - Small Outline Dual In-line Memory Module (SODIMM)
  - Dual In-line Memory Module (DIMM)
- Double Data Rate (DDR) iterations
- Error-correcting code (ECC) vs. non-ECC RAM
- Channel configurations

✓ **3.5 Given a scenario, install and configure motherboards, central processing units (CPUs), and add-on cards.**

- Motherboard form factors
  - Advanced Technology Extended (ATX)
  - microATX
  - Information Technology eXtended (ITX)
- Motherboard connector types
  - Peripheral Component Interconnect (PCI)
  - PCIe
  - Power connectors
  - SATA
  - eSATA

- Headers
- M.2

- Motherboard compatibility
  - CPU socket types
    - Advanced Micro Devices, Inc. (AMD)
    - Intel
  - Multisocket
- BIOS/Unified Extensible Firmware Interface (UEFI) settings
  - Boot options
  - USB permissions
  - Trusted Platform Module (TPM) security features
  - Fan considerations
  - Secure Boot
  - Boot password
  - BIOS password
  - Temperature monitoring
- Virtualization support
- Encryption
  - TPM
  - Hardware security module (HSM)
- CPU architecture
  - x86/x64
  - Advanced RISC [Reduced Instruction Set Computer] Machine (ARM)
  - Core configurations
- Cooling
  - Fans
  - Heat sink
  - Thermal paste/pads
  - Liquid

The computers we use daily, from the largest servers to the smallest smart watches and everything in between, are collections of different electronic components and software working together in a system. Digital computers have been around since the late 1930s, so they aren't new news. As you would imagine, though, their looks and functionality have evolved considerably since then.

As technology improved over the years, computers got smaller and faster, and inventors added features that required new hardware devices. Inspired by typewriters, keyboards were added for ease of input. Visual displays required a monitor, a video card, and a standard interface between them. Sound was provided by a different expansion card and speakers. Because of the newly added features, PCs were modular by necessity. That is, if a new feature or functionality were needed, a new component could be added. Or if a part failed, it could be replaced by a new one. Along the way in the late 1960s, the term *personal computer (PC)* was coined to differentiate between computers designed to be used by one person versus other options, such as a mainframe or one where multiple users share a processor. This book, and the CompTIA A+ exams, focus on PC hardware and software.

Unless specifically mentioned otherwise, the terms *PC* and *computer* are used interchangeably throughout this book.

Much of the computing industry today is focused on smaller devices, such as laptops, tablets, smartphones, and Internet of Things (IoT) devices. Laptops have outsold desktop computers since 2005, and it seems that everyone has a smartphone glued to their hand today. IoT devices (such as smart home automation devices and virtual assistants like Amazon Echo and Google Nest, which we will cover more in depth in Volume 1 Chapter 8, "Network Services,") have exploded in popularity from about 10 billion in 2019 to an expected 40 billion by 2030. Smaller devices require the same components as do their bigger desktop-sized cousins. Of course, the components are smaller, and by necessity most of the time they are integrated into the same circuit board. The functionality of the individual parts is still critical, though, so what you learn here will serve you well regardless of the type of device you're working on.

Even though all parts inside a computer case are important, some are more important than others. You'll sometimes hear people refer to the "big three" computer parts, which are the motherboard, processor, and memory. Without these, a computer won't work, whereas if a sound card fails, it's probably not going to make the entire system inoperable. In this chapter, you will learn how to identify, install, and configure the big three, which are critical

to computer functionality. You'll also learn about cooling systems, because too much heat will cause components to melt, which could make for a very bad day for the computer user.

With small computing devices being far more prevalent, you might wonder why we are starting the book talking about components that are primarily in desktop and larger systems. We intentionally cover this material first because it's important to understand what each component does individually. It helps you isolate the functionality of each piece of the system, which is critical for troubleshooting and fixing problems. Besides, if you are working (or want to be working) as a computer technician, you will likely see plenty of desktop-sized PCs in your travels!

Understanding Motherboards

The spine of the computer is the *motherboard*, otherwise known as the *system board* or *mainboard*. This is the *printed circuit board (PCB)*, which is a conductive series of pathways laminated to a nonconductive substrate that lines the bottom of the computer and is often of a uniform color, such as green, brown, blue, black, or red. It is the most important component in the computer because it connects all the other components together. Figure 1.1 shows a typical PC system board, as seen from above. All other components are attached to this circuit board. On the system board, you will find the central processing unit (CPU) slot or integrated CPU, underlying circuitry, expansion slots, video components, random access memory (RAM) slots, and a variety of other chips. We will be discussing each of these components throughout this book.

Motherboard Form Factors

You can find hundreds, if not thousands, of different motherboards on the market today. It can be overwhelming trying to figure out which one is needed. To help, think of motherboards in terms of which type they are, based on a standard classification. System boards are classified by their *form factor* (the design of the component), such as ATX and ITX. Exercise care and vigilance when acquiring a motherboard and components separately. Motherboards will have different expansion slots, support certain processors and memory, and fit into some cases but not others. Be sure that the other parts are physically compatible with the motherboard you choose.

The CompTIA A+ exam 220-1201 objective 3.5, "Given a scenario, install and configure motherboards, central processing units (CPUs), and add-on cards," specifically calls out ATX, microATX, and ITX motherboard form factors. Be sure to understand the differences between the three standards.

FIGURE 1.1 A typical motherboard

Advanced Technology Extended

Intel developed the *Advanced Technology Extended (ATX)* motherboard in the mid-1990s to improve on the classic AT-style motherboard architecture that had ruled the PC world for many years. The ATX motherboard has the processor and memory slots at right angles to the expansion cards, like the one in Figure 1.1. This arrangement puts the processor and memory in line with the fan output of the power supply, allowing the processor to run cooler. And because those components are not in line with the expansion cards, you can

install full-length expansion cards—adapters that extend the full length of the inside of a standard computer case—in an ATX motherboard machine. ATX (and its derivatives, such as microATX) is the primary PC motherboard form factor in use today. Standard ATX motherboards measure 12″ × 9.6″ (305mm × 244mm). A slightly smaller version of an ATX motherboard is *microATX*. These boards are 9.6″ (244mm) square, and because of their smaller size typically have fewer connectors for plugging in expansion cards and peripherals. ATX and microATX motherboards use the same case mounting points, so a microATX board will fit into a standard ATX case. The opposite might not be true, however, depending on case dimensions.

We will discuss expansion cards in more depth (and cover that part of Exam Objective 3.5) in Volume 1 Chapter 2, "Expansion Cards, Storage Devices, and Power Supplies."

Information Technology Extended

The *Information Technology Extended (ITX)* line of motherboard form factors was developed by VIA Technologies in the early 2000s as a low-power, small form factor (SFF) board for specialty uses, including home-theater systems, compact desktop systems, gaming systems, and embedded components. ITX itself is not an actual form factor but a family of form factors. The family consists of the following form factors:

- **Mini-ITX:** 6.7″ × 6.7″ (170mm × 170mm)
- **Nano-ITX:** 4.7″ × 4.7″ (120mm × 120mm)
- **Pico-ITX:** 3.9″ × 2.8″ (100mm × 72mm)
- **Mobile-ITX:** 2.4″ × 2.4″ (60mm × 60mm)

The mini-ITX motherboard has four mounting holes that line up with three or four of the holes in the ATX and micro-ATX form factors. In mini-ITX boards, the rear interfaces are placed in the same location as those on the ATX motherboards. These features make mini-ITX boards compatible with ATX cases. This is where the mounting compatibility ends, because despite the PC compatibility of the other ITX form factors, they are used in embedded systems, such as set-top boxes, home entertainment systems, and smartphones, and lack the requisite mounting and interface specifications. Figure 1.2 shows the three larger forms of ITX motherboards, next to two ATX motherboards for comparison.

System Board Components

Now that you understand the basic types of motherboards and their form factors, it's time to look at the key characteristics and components of the motherboard and, where applicable, their locations relative to each other. The following list summarizes key concepts you need to know about motherboards:

- Bus architecture
- Chipsets

- Connector types (expansion slots)
- Memory slots and cache
- Central processing units and processor sockets
- Power connectors
- Onboard nonvolatile storage connectors (such as for hard drives)
- Motherboard headers
- BIOS/UEFI/firmware
- CMOS and CMOS battery
- Front-panel connectors and headers

FIGURE 1.2 ITX motherboards

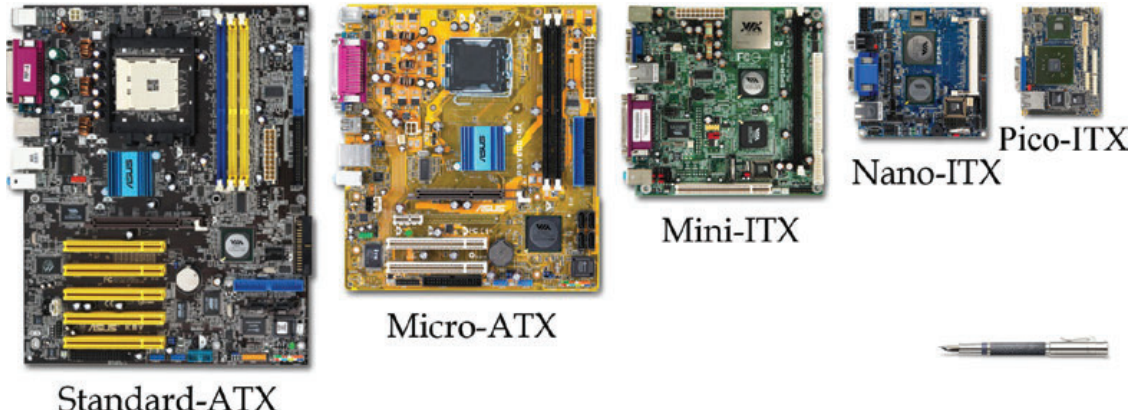

Source: VIA Mini-ITX Form Factor Comparison / VIA Gallery / Wikimedia Commons / CC BY 2.0

In the following sections, you will learn about some of the most common components of a motherboard, what they do, and where they are located on the motherboard. We'll show what each component looks like so that you can identify it on most any motherboard that you run across. In the case of some components, this chapter provides only a brief introduction, with more detail to come in later chapters.

Bus Architecture

In a PC, data is sent from one component to another via a *bus*, which is a common collection of signal pathways. In the very early days, PCs used serial buses, which sent one bit at a time and were painfully slow. Brilliant engineers realized that they could redesign the bus and send 8 bits at a time (over synchronized separate lines), which resulted in a big speed increase. This was known as a *parallel bus*.

The downside of parallel communications is the loss of circuit length (how long the circuit could be) and throughput (how much data could move at one time). The signal could travel only a short distance, and the amount of data was limited due to the careful synchronization needed between separate lines, the speed of which must be controlled to limit skewing the arrival of the individual signals at the receiving end.

What was once old is new again, as engineers have discovered methods to make serial transmissions work at data rates that are many times faster than parallel signals. Therefore, nearly everything you see today uses a *serial bus*. The only limitation of serial circuits is in the capability of the transceivers, which tends to grow over time due to technical advancements. Examples of specifications that have heralded the dominance of serial communications are Serial Advanced Technology Attachment (Serial ATA, or SATA), Universal Serial Bus (USB), and Peripheral Component Interconnect Express (PCIe).

The term *bus* is also used in any parallel or bit-serial wiring implementation where multiple devices can be attached at the same time in parallel or in series (daisy-chained). Examples include Small Computer System Interface (SCSI), USB, and Ethernet.

On a motherboard, several different buses are used. Motherboard connector types of various architectures, such as PCIe, are included to allow for the insertion of external devices or adapters. Other types of buses exist within the system to allow communication between the CPU, RAM, and other components with which data must be exchanged. Except for CPU sockets and memory slots, there are no insertion points for devices in many closed bus systems because no expansion cards exist for those features.

The various buses throughout a given computer system can be rated by their bus speeds. The higher the bus speed, the higher its performance. In some cases, various buses must be synchronized for proper performance, such as the system bus and any expansion buses that run at the front-side bus speed. Other times, one bus will reference another for its own speed. The internal bus speed of a CPU is derived from the front-side bus clock, for instance. The buses presented throughout this chapter are accompanied by their speeds, where appropriate.

Chipsets

A *chipset* is a collection of chips or circuits that perform interface and peripheral functions for the processor. This collection of chips is usually the circuitry that provides interfaces for memory, expansion cards, and onboard peripherals, and it generally dictates how a motherboard will communicate with the installed peripherals.

Chipsets are usually given a name and model number by the original manufacturer. For example, B550 and X570 are chipsets that support Advanced Micro Devices, Inc. (AMD) processors, and Z490 and H410 are Intel motherboard chipsets. Typically, the manufacturer and model tell you that your particular chipset has a certain set of features, such as:

- Type of CPU and RAM supported
- Onboard video capabilities

- Motherboard expansion slot types, versions, and speeds
- Number of input/output ports and versions
- Support for BIOS/UEFI security

Don't worry if you don't understand what those features are just yet—we'll cover them throughout this chapter. Also don't worry about memorizing any chipset names—you can look them up online to understand their features. For now, know that the motherboard chipset and CPU must be designed to work together.

Motherboard Connector Types (Expansion Slots)

The most visible parts of any motherboard are the connectors for internal expansion cards, also called *expansion slots*. These are small plastic slots, usually from 1 to 6 inches long and approximately 1/2 inch wide. As their name suggests, these slots are used to install various devices in the computer to expand its capabilities. Some expansion devices that might be installed in these slots include video, network, sound, and disk interface cards.

If you look at the motherboard in your computer, you will more than likely see one of the main types of expansion slots used in computers today, which are PCI and PCIe. The exam objectives mention several other motherboard connector types as well such as SATA and M.2. We will cover all of them throughout this chapter, but for now, we will focus on how to visually identify the different expansion slots on the motherboard.

Exam Objective 3.5 lists several motherboard connector types you should be familiar with. They include PCI and PCIe.

PCI Expansion Slots

It's now considered an old technology, but many motherboards in use today still contain 32-bit *Peripheral Component Interconnect (PCI)* slots (32-bit means that data is handled 32 bits at a time). They are easily recognizable because they are only around 3″ long and classically white, although modern boards take liberties with the color. PCI slots became extremely popular with the advent of Pentium-class processors in the mid-1990s. Although popularity has shifted from PCI to PCIe, the PCI slot's service to the industry cannot be ignored; it has been an incredibly prolific architecture for many years.

PCI expansion buses operate at 33 MHz or 66 MHz (version 2.1) over a 32-bit (4-byte) channel, resulting in data rates of 133 MBps and 266 MBps, respectively, with 133 MBps being the most common, server architectures excluded. PCI is a shared-bus topology, however, so mixing 33 MHz and 66 MHz adapters in a 66 MHz system will slow all adapters to 33 MHz. Older servers might have featured 64-bit PCI slots as well, since version 1.0, which double the 32-bit data rates. See the sidebar "Arriving at the Exact Answer" for help with understanding the math involved in frequencies and bit rates.

Arriving at the Exact Answer

To get the math exactly right when dealing with frequencies and data rates ending in 33 and 66, you have to realize that every 33 has an associated one-third (1/3), and every 66 has an associated two-thirds (2/3). The extra quantities are left off of the final result but must be added back on to get the math exactly right. The good news is that omitting these small values from the equation still gets you close, and a bit of experience with the numbers leads to being able to make the connection on the fly.

PCI slots and adapters are manufactured in 3.3V and 5V versions. Universal adapters are keyed to fit in slots based on either of the two voltages. The notch in the card edge of the common 5V slots and adapters is oriented toward the front of the motherboard and the notch in the 3.3V adapters toward the rear. Figure 1.3 shows several PCI expansion slots. Note the shorter 5V 32-bit slot in the foreground and the longer 3.3V 64-bit slots. Also notice that a universal 32-bit card, which has notches in both positions, is inserted into and operates fine in the 64-bit 3.3V slot in the background.

FIGURE 1.3 PCI expansion slots

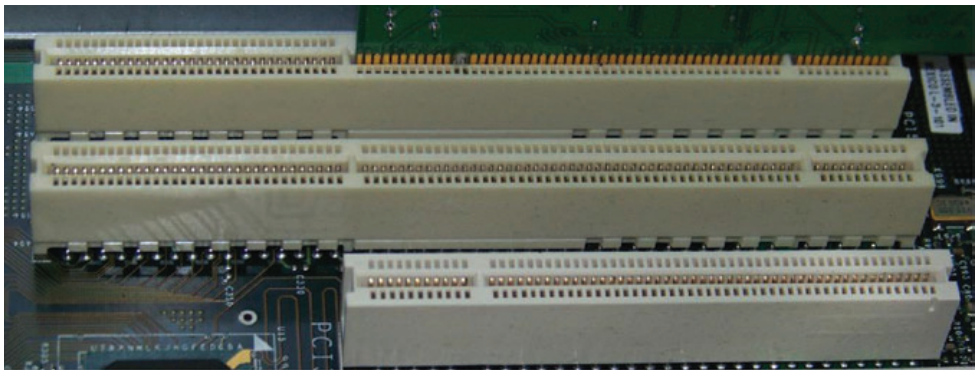

PCIe Expansion Slots

The most common expansion slot architecture used on motherboards is *PCI Express (PCIe)*. It was designed to be a replacement for PCI, as well as an older video card standard called accelerated graphics port (AGP). PCIe has the advantage of being faster than AGP while maintaining the flexibility of PCI. PCIe has no plug compatibility with either AGP or PCI. Some modern PCIe motherboards can be found with regular PCI slots for backward compatibility, but AGP slots have not been included for many years.

AGP was a video card–only standard introduced in 1997 and became obsolete in the late 2000s. If you see an old motherboard with a brown expansion slot parallel to white PCI slots and set back a bit from the rear of the case, it's probably AGP.

PCIe is casually referred to as a bus architecture to simplify its comparison with other bus technologies. True expansion buses share total bandwidth among all slots, each of which taps into different points (aka multipoint) along the common bus lines. In contrast, PCIe uses a switching component with point-to-point connections to slots, which gives each component full use of the corresponding bandwidth and making connections faster. Furthermore, unlike other expansion buses, which have parallel architectures, PCIe is a serial technology, striping data packets across multiple serial paths to achieve higher data rates.

PCIe uses the concept of *lanes*, which are the switched point-to-point signal paths between any two PCIe components. Each lane that the switch interconnects between any two intercommunicating devices comprises a separate pair of wires for both directions of traffic. Each PCIe pairing between cards requires a negotiation for the highest mutually supported number of lanes. The single lane, or combined collection of lanes that the switch interconnects between devices, is referred to as a *link*.

There are seven different link widths supported by PCIe, designated x1 (pronounced "by 1"), x2, x4, x8, x12, x16, and x32, with x1, x4, and x16 being the most common. A slot that supports a particular link width is of a physical size related to that width because the width is based on the number of lanes supported, requiring a related number of wires. As a result, a x8 slot is longer than a x1 slot but shorter than a x16 slot. Every PCIe slot has a 22-pin portion in common toward the rear of the motherboard, which you can see in Figure 1.4, in which the rear of the motherboard is to the left. (Figure 1.4 shows, from top to bottom, PCIe x4, PCIe x16, PCIe x1, PCIe x16, and PCI slots.) These 22 pins comprise mostly voltage and ground leads.

FIGURE 1.4 PCIe expansion slots

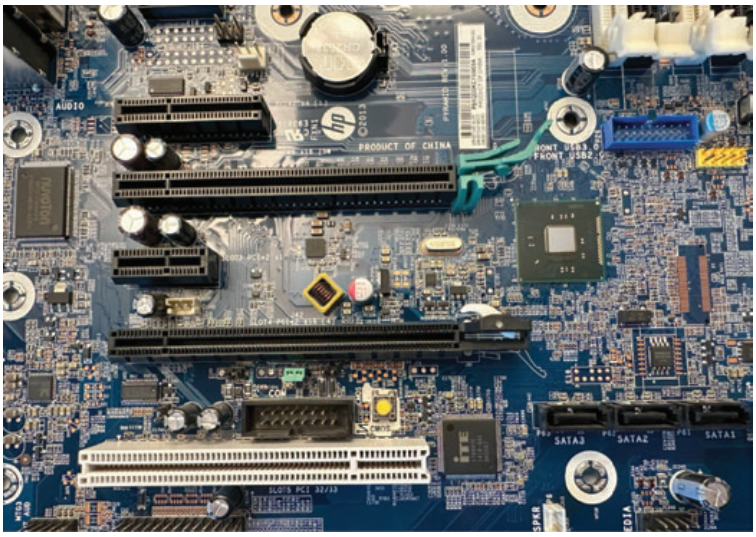

In addition to PCIe link widths, you should be aware of PCIe standards. The oldest version, 1.X, was released in 2003, running at 250 MBps per direction, or 500 Mbps bidirectionally. Use the doubled rate when comparing PCIe to other expansion buses because

those other rates are for bidirectional communication. This means that the original, slowest 500 MBps bidirectional link of a x1 slot in the first version of PCIe was comparable to PCI's best, a 64-bit slot running at 66 MHz and producing a throughput of 533 MBps.

Bidirectional means that data flows in both directions, often simultaneously. *Unidirectional* means data flows in only one direction.

Combining lanes simply results in a linear multiplication of these rates. For example, a PCIe 1.1 x16 slot is capable of 4 GBps of throughput in each direction, which is 16 times the 250 MBps x1 rate. Bidirectionally, this fairly common slot produces a throughput of 8 GBps. Each subsequent PCIe specification doubles this throughput. Following this math through the versions, PCIe 6.0 has a single-lane data rate of 8 GBps, for a total x16 bidirectional bandwidth of 256 GBps. It's really fast.

Real World Scenario

The Speed of Innovation

Newer standards always offer faster data transfer rates and greater throughput. Usually when the standard is released, people shake their heads at how mind-bogglingly fast they are, and then the cycle repeats itself a few years later. Hardware devices that support the standard are often one or two years behind the release of the standard.

For example, PCIe 4.0 was formally ratified by the PCI Special Interest Group (PCI-SIG) in October 2017, seven years after PCIe 3.0. It doubled the data rates of the previous standard and provided good performance for virtual and augmented reality applications.

PCIe 5.0 was released in May 2019 and of course doubled up on PCIe 4.0. This helped support machine learning and artificial intelligence, as well as high-speed data transfers needed in data centers.

Continuing the journey is PCIe 6.0, released in January 2022. Devices supporting this standard started to hit the market in 2024. And PCIe 7.0 has a 2025 release date and devices are expected in 2027.

The key thing to remember about PCIe versions is that each generation doubles the speed before it. So if you remember that PCIe 1.0 x1 was 500 Mbps, you can use a calculator to figure out any combination of standard and link width. If you want to read further on PCIe details, visit `https://pcisig.com/specifications`.

The good news is that you don't have to worry about compatibility between cards and slots of different standards. A PCIe 5.0 x1 card will fit into a PCIe 4.0 x1 slot and vice versa. The device will simply operate at the speed of the slower standard.

Refer back to Figure 1.4, which shows two PCIe x16 slots. It might be a bit hard to see in the picture, but on the right side of the slots there is a latch and tab that secures the x16 adapter in place by its hooks. Any movement of these high-performance adapters can result in temporary failure or poor performance, so the latch is critical. It can also make removing the card a little tricky if the latch isn't released first.

Using Riser Cards

Most PC expansion cards plug directly into the motherboard. A special type of expansion card, called a *riser card*, provides additional slots that other expansion cards plug into. The other expansion cards are then parallel to the motherboard, as opposed to perpendicular. Figure 1.5 shows an example of a riser card. It's an older example, but it illustrates the concept well.

Riser cards aren't often found in desktop PCs today but do still find some use in rack-mounted servers with low-profile cases. The motherboard must be designed to accept a riser card, and the case needs to be built for it as well, for the external portion of the expansion card to be accessible.

FIGURE 1.5 Riser card in a motherboard

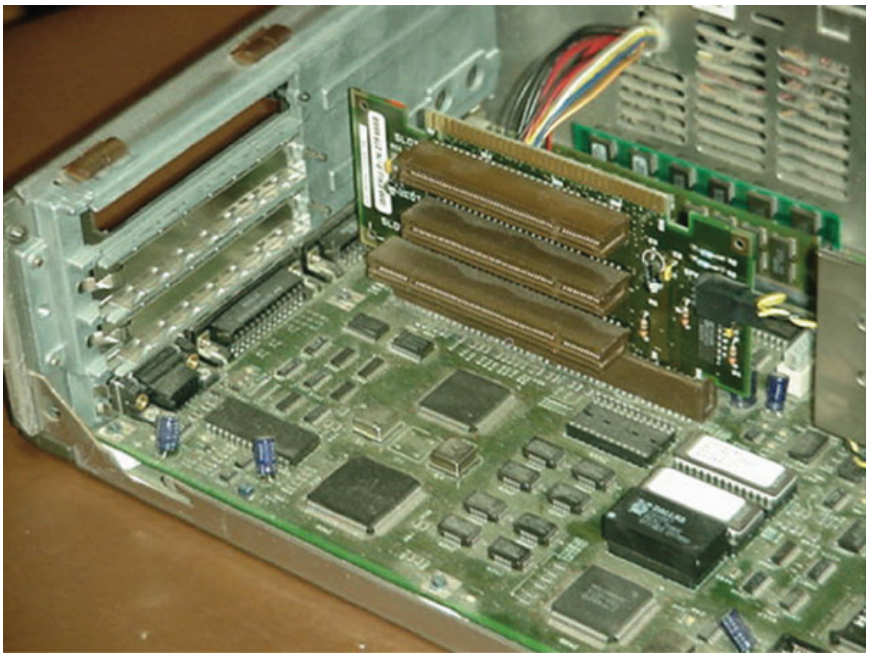

Source: Dale Mahalko / Wikimedia Commons / CC BY SA 3.0

Memory Slots and Cache

Memory, or *random access memory (RAM)*, slots are the next most notable slots on a motherboard. These slots are designed for the modules that hold memory chips that make up primary memory, which is used to store currently used data and instructions for the CPU. Many types of memory are available for PCs today. In this chapter, you will become familiar with the appearance and specifications of the slots on the motherboard so that you can identify them and appropriately install or replace RAM.

For the most part, PCs today use memory chips arranged on a small circuit board. A *dual in-line memory module (DIMM)* is one type of circuit board. Today's DIMMs differ in the number of conductors, or pins, that each particular physical form factor uses. Some common examples include 168-, 184-, 240-, and 288-pin configurations. In addition, laptop memory comes in smaller form factors known as small outline DIMMs (SODIMMs) and Micro-DIMMs. More detail on memory packaging and the technologies that use them can be found later in this chapter in the section "Understanding Memory."

Memory slots are easy to identify on a motherboard. Classic DIMM slots were usually black and, like all memory slots, were placed very close together. DIMM slots with color-coding are more common these days, however. The color-coding of the slots acts as a guide to the installer of the memory. See the section "Channel Configurations," later in this chapter, for more on the purpose of this color-coding. Consult the motherboard's documentation to determine the specific modules allowed as well as their required orientation. The number of memory slots varies from motherboard to motherboard, but the structure of the different slots is similar. Metal pins in the bottom make contact with the metallic pins on each memory module. Small metal or plastic tabs on each side of the slot keep the memory module securely in its slot. Figure 1.6 shows four memory slots, with the CPU socket included for reference.

Another type of memory common in PCs is *cache memory*, which is small and fast and logically sits between the CPU and RAM. Cache is a very fast form of memory forged from static RAM, which is discussed in detail in the section "Understanding Memory," later in this chapter. Cache improves system performance by predicting what the CPU will ask for next and prefetching this information before being asked. This paradigm allows the cache to be smaller in size than the RAM itself. Only the most recently used data and code or that which is expected to be used next is stored in cache.

You'll see three different cache designations:

Level 1 Cache *L1 cache* is the smallest and fastest, and it's on the processor die itself. In other words, it's an integrated part of the manufacturing pattern that's used to stamp the processor pathways into the silicon chip. You can't get any closer to the processor than that.

Though the definition of L1 cache has not changed much over the years, the same is not true for other cache levels. L2 and L3 cache used to be on the motherboard but now have moved on-die in most processors as well. The biggest differences are the speed and whether they are shared.

Level 2 Cache *L2 cache* is larger but a little slower than L1 cache. For processors with multiple cores, each core will generally have its own dedicated L1 and L2 caches. A few processors share a common L2 cache among the cores.

Level 3 Cache *L3 cache* is larger and slower than L1 or L2, and is usually shared among all processor cores.

FIGURE 1.6 Double data rate (DDR) memory slots

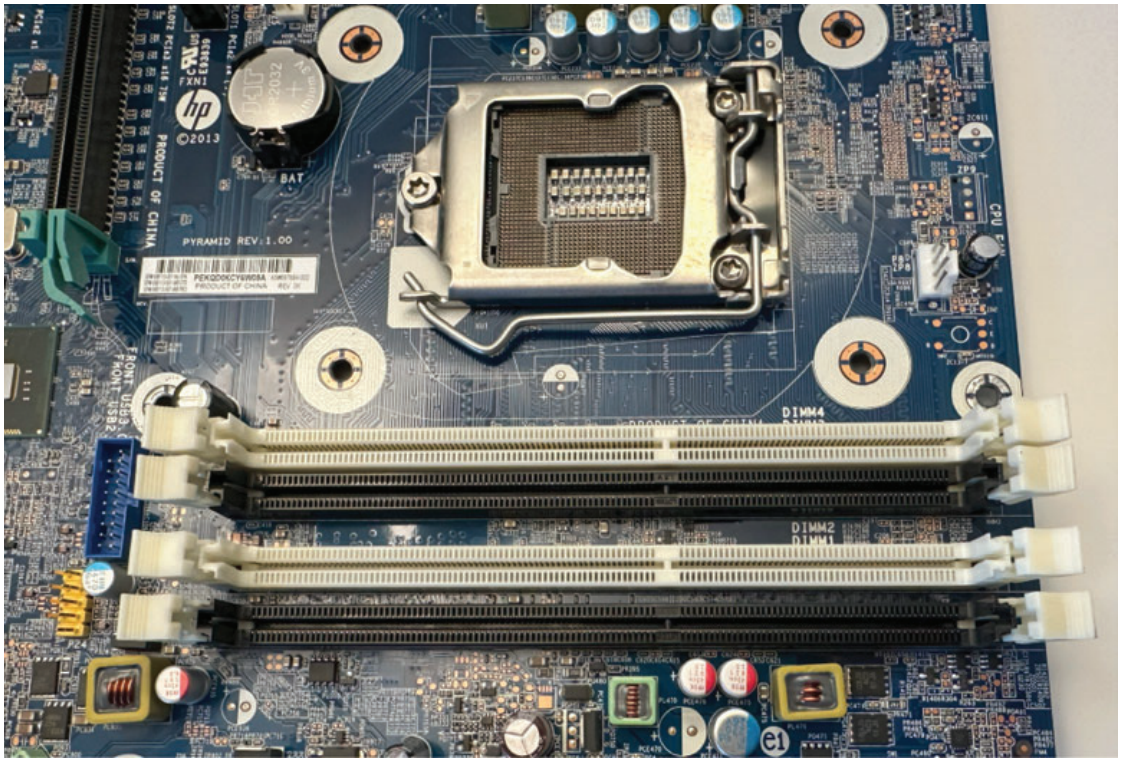

The typical increasing order of capacity and distance from the processor die is L1 cache, L2 cache, L3 cache, RAM, and HDD/SSD (hard disk drive and solid-state drive—more on these in Volume 1 Chapter 2). This is also the typical decreasing order of speed. The following list includes representative capacities of these memory types. The cache capacities are for each core of the 10th generation Intel Core i7 processor. The other capacities are simply modern examples.

- **L1 cache:** 80 KB (32 KB for instructions and 48 KB for data)
- **L2 cache:** 512 KB
- **L3 cache:** 8–16 MB
- **RAM:** 16–256 GB
- **HDD/SSD:** 100s of GB to several TB

One way to find out how much cache your system has is to use a utility such as CPU-Z, as shown in Figure 1.7. CPU-Z is freeware that can show you the amount of cache, processor name and number, motherboard and chipset, and memory specifications. It can be found at `www.cpuid.com`.

FIGURE 1.7 Cache in a system

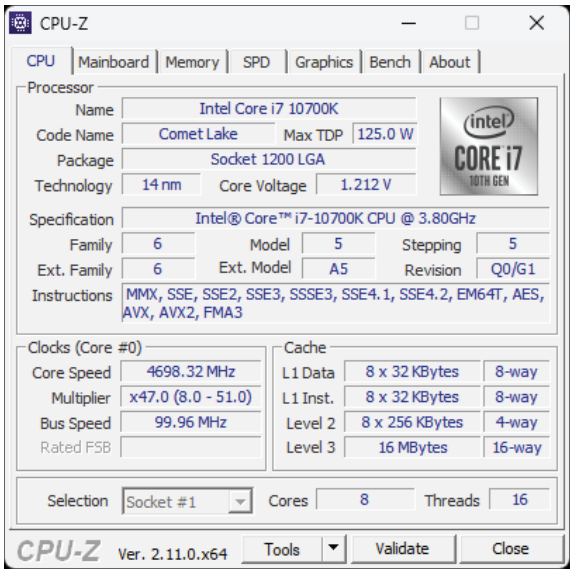

Figure 1.7 shows L1 Data and L1 Inst. caches. Many CPUs will split L1 cache into cache for data (L1D) or instructions (L1I). It's highly unlikely you'll be tested on this, but it's interesting to know.

Central Processing Unit and Processor Socket

The "brain" of any computer is the *central processing unit (CPU)*. There's no computer without a CPU. There are many different types of processors for computers—so many, in fact, that you will learn about them later in this chapter in the section "Understanding Processors." Typically, in today's computers, the processor is the easiest component to identify on the motherboard. It is usually the component that has either a fan or a heat sink (usually both) attached to it (as shown in Figure 1.8). These devices are used to draw away and disperse the heat that a processor generates. This is done because heat is the enemy of microelectronics. Today's processors generate enough heat that, without heat dispersal, they would permanently damage themselves and the motherboard in a matter of minutes, if not seconds.

CPU sockets are almost as varied as the processors that they hold. Sockets are basically flat and have several columns and rows of holes or pins arranged in a square, as shown in Figure 1.9. The left socket is known as Socket AM4, made for AMD processors such as

the Ryzen, and has holes to receive the pins on the CPU. This is known as a *pin grid array (PGA)* arrangement for a CPU socket. The holes and pins are in a row/column orientation, an array of pins. The right socket is known as LGA 1200, and there are spring-loaded pins in the socket and a grid of lands on the CPU. The *land grid array (LGA)* is a newer technology that places the delicate pins (yet more sturdy than those on chips) on the cheaper motherboard instead of on the more expensive CPU, opposite to the way that the aging PGA does. The device with the pins has to be replaced if the pins become too damaged to function. The PGA and LGA are mentioned again later in this chapter in the section "Understanding Processors."

FIGURE 1.8 Two heat sinks, one with a fan

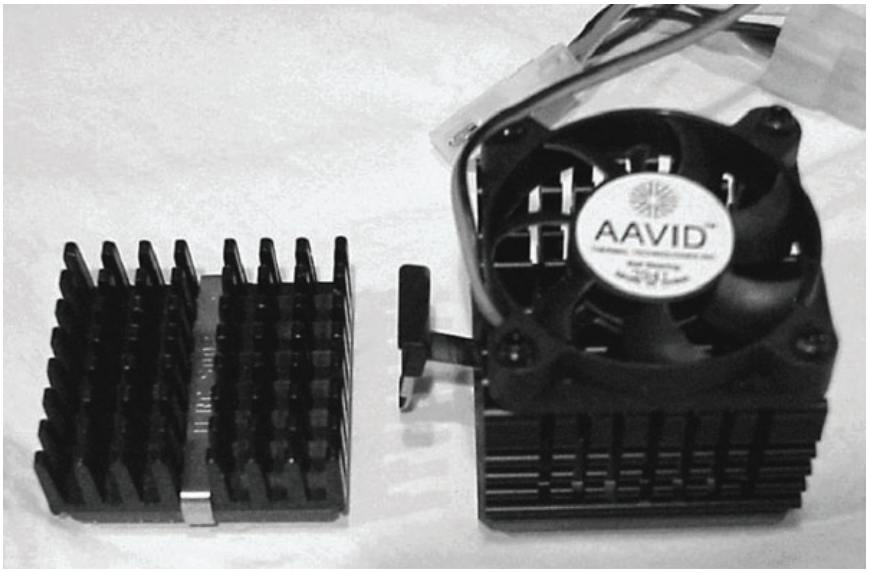

FIGURE 1.9 CPU socket examples

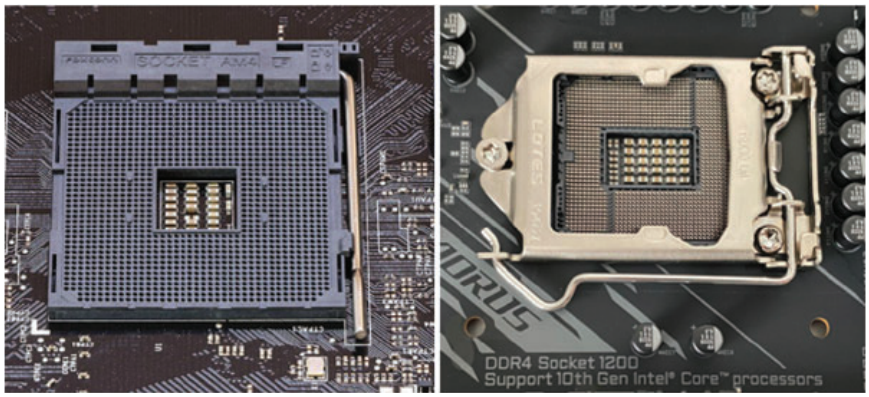

Modern CPU sockets have a mechanism in place that reduces the need to apply considerable force to the CPU to install a processor, which was necessary in the early days of personal computing. Given the extra surface area on today's processors, excessive pressure applied in the wrong manner could damage the CPU packaging, its pins, or the motherboard itself. For CPUs based on the PGA concept, *zero insertion force (ZIF)* sockets are exceedingly popular. ZIF sockets use a plastic or metal lever on one of the two lateral edges to lock or release the mechanism that secures the CPU's pins in the socket. The CPU rides on the mobile top portion of the socket, and the socket's contacts that mate with the CPU's pins are in the fixed bottom portion of the socket. The image of Socket AM4 shown on the left in Figure 1.9 illustrates the ZIF locking mechanism at the right edge of the socket.

For processors based on the LGA concept, a socket with a different locking mechanism is used. Because there are no receptacles in either the motherboard or the CPU, there is no opportunity for a locking mechanism that holds the component with the pins in place. LGA-compatible sockets, as they're called despite the misnomer, have a lid that closes over the CPU and is locked in place by an L-shaped arm that borders two of the socket's edges. The nonlocking leg of the arm has a bend in the middle that latches the lid closed when the other leg of the arm is secured. The right image in Figure 1.9 shows an LGA socket with no CPU installed and the locking arm secured over the lid's tab, along the bottom edge.

Listing out all the desktop PC socket types you might encounter would take a long time. Instead, we'll give you a sampling of some that you might see. The first thing you might notice is that sockets are made for Intel or AMD processors, but not both. Keep that compatibility in mind when replacing a motherboard or a processor. Make sure that the processor and motherboard were designed for each other (even within the Intel or AMD families); otherwise, they won't fit each other and won't work. Table 1.1 lists some common desktop socket/CPU relationships. Servers and laptops/tablets generally have different sockets altogether, although some CPU sockets will support processors designed for desktops or servers.

TABLE 1.1 Desktop PC socket types and the processors they support

Socket	Released	Type	Pin count	Processors
Socket sTR5	2023	LGA	4844	AMZ Ryzen Threadripper
LGA 1700	2021	LGA	1700	Intel Alder Lake and Raptor Lake
LGA 1200	2020	LGA	1200	Intel Comet Lake and Rocket Lake
Socket AM4	2017	PGA	1331	AMD Ryzen 3, Ryzen 5, Ryzen 7, Ryzen 9, Athlon 200GE
LGA 2066 (Socket R4)	2017	LGA	2066	Intel Skylake-X and Kaby Lake-X
LGA 1151 (Socket H4)	2015	LGA	1151	Intel Skylake, Kaby Lake, and Coffee Lake

Socket	Released	Type	Pin count	Processors
Socket FM2+	2014	PGA	906	AMD Kaveri and Godavari
Socket AM1	2014	PGA	721	AMD Athlon and Sempron

Real World Scenario

Playing the Name Game

Back in simpler times, Intel marketed only one or two processors at a time. For example, when the 80486DX was on the market, they didn't have much else. Then the Pentium came along, followed by the Celeron (which everyone knew was a stripped-down Pentium), and the Pentium M, which was the mobile version of the Pentium. The active roster of processors started to expand.

Now, there are a lot of active chips, including the Xeon and the Core series (i9, i7, i5, and i3), and even the Pentium and Celeron are still kicking around. And with new innovation coming nearly every year, it's difficult to keep track of what's current.

In about 2013, Intel began to publicize its project code names a bit more to make it easier to keep track of the new technology bundles. No longer were techs able to remember that an i7 required a specific type of socket, but they needed to know *which* i7 they were dealing with. For example, a Comet Lake (2020) i7 requires an LGA 1200 socket, as does the Rocket Lake, whereas the Alder Lake (2021) i7s use an LGA 1700.

You aren't going to need to memorize all possible Intel chip and motherboard socket combinations for the exam or the real world. In the real world, the Internet's always there to remind you of which socket you need for a specific processor. To give you a flavor, though, Table 1.2 covers a few of Intel's recent CPU architecture codenames and the associated processor model numbers.

TABLE 1.2 Select Intel desktop processors

Name (Year)	Gen	Socket	Core i9	Core i7	Core i5	Core i3
Raptor Lake-S Refresh (2023)	14th	LGA1700	149xx	147xx	146xx 145xx 144xx	141xx

TABLE 1.2 Select Intel desktop processors *(continued)*

Name (Year)	Gen	Socket	Core i9	Core i7	Core i5	Core i3
Raptor Lake (2022)	13th	LGA 1700	139xx	137xx	136xx 135xx 134xx	131xx
Alder Lake (2021)	12th	LGA 1700	129xx	127xx	126xx	n/a
Rocket Lake (2020)	11th	LGA 1200	119xx	117xx	116xx 115xx 114xx	n/a

To make things even more interesting, you might hear references to the *n*th generation of the i5, i7, or i9 processor. For example, Raptor Lake is considered 13th generation. Intel codenames over the past several years have had some variation of a "Lake" name, such as Coffee Lake (2017), Comet Lake (2019), and Rocket Lake (2020). Perhaps the names keep it interesting for the engineers.

When you go to the parts store, though, you're probably not going to ask for a Rocket Lake i7. You'll ask for an i7, see what they have in stock—usually they will have the current generation and possibly the previous one, and then figure out what motherboard you need. (Or alternatively, based on your motherboard, you'll figure out which i7 you can use.) For example, the motherboard may advertise that it supports "12th and 13th generation" Intel processors. If your processor number starts with a 12xxx or a 13xxx, you should be good to go.

What does all of this mean you need to do? First, get comfortable knowing that a new architecture may require a different motherboard. Second, feel free to memorize some of the more common combinations of motherboards and processors if you find it helpful, although you won't be tested on it. It may be handy in the real world to know that if you have an 11th gen Intel chip, you need a board with an LGA 1200 socket. (Or you can always Google it!)

For A+ Exam Objective 3.5, be familiar with compatibility of CPU socket types (AMD and Intel) as well as multisocket motherboards.

Multisocket Motherboards

When it comes to motherboard compatibility, the two biggest things to keep in mind are the processor type and the case. If either of those are misaligned with what the motherboard supports, you're going to have problems.

Thus far, as we've talked about desktop motherboards and their CPU sockets, we have shown examples of boards that have just one socket. There are motherboards that have more than one CPU socket and conveniently, they are called *multisocket* (typically written as two words) motherboards. Figure 1.10 shows a two-socket motherboard made by GIGABYTE. The two CPU sockets are easily identifiable and note that each CPU socket has eight dedicated memory slots. Multisocket motherboards will generally use the same CPU sockets that single-socket motherboards use.

FIGURE 1.10 GIGABYTE multisocket motherboard

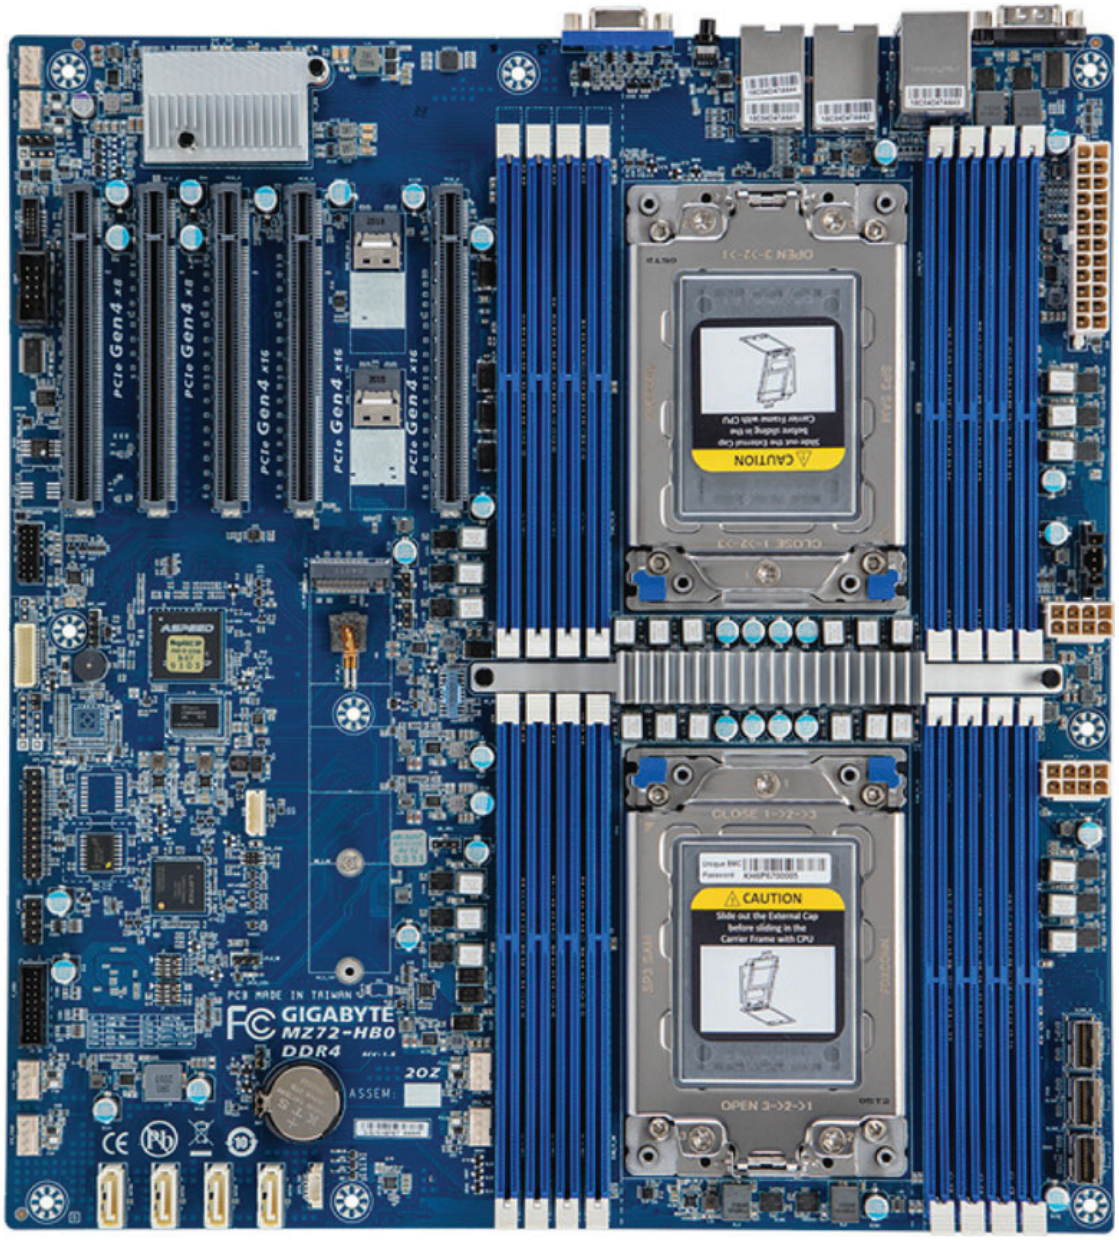

Mobile Motherboards

In small mobile devices, space is at a premium. Some manufacturers will use standard small-factor motherboards, but most create their own boards to fit inside specific cases. An example of an oddly shaped Dell laptop motherboard is shown in Figure 1.11. When replacing a laptop motherboard, you almost always need to use one from the exact same model; otherwise it won't fit inside the case.

FIGURE 1.11 Dell laptop motherboard

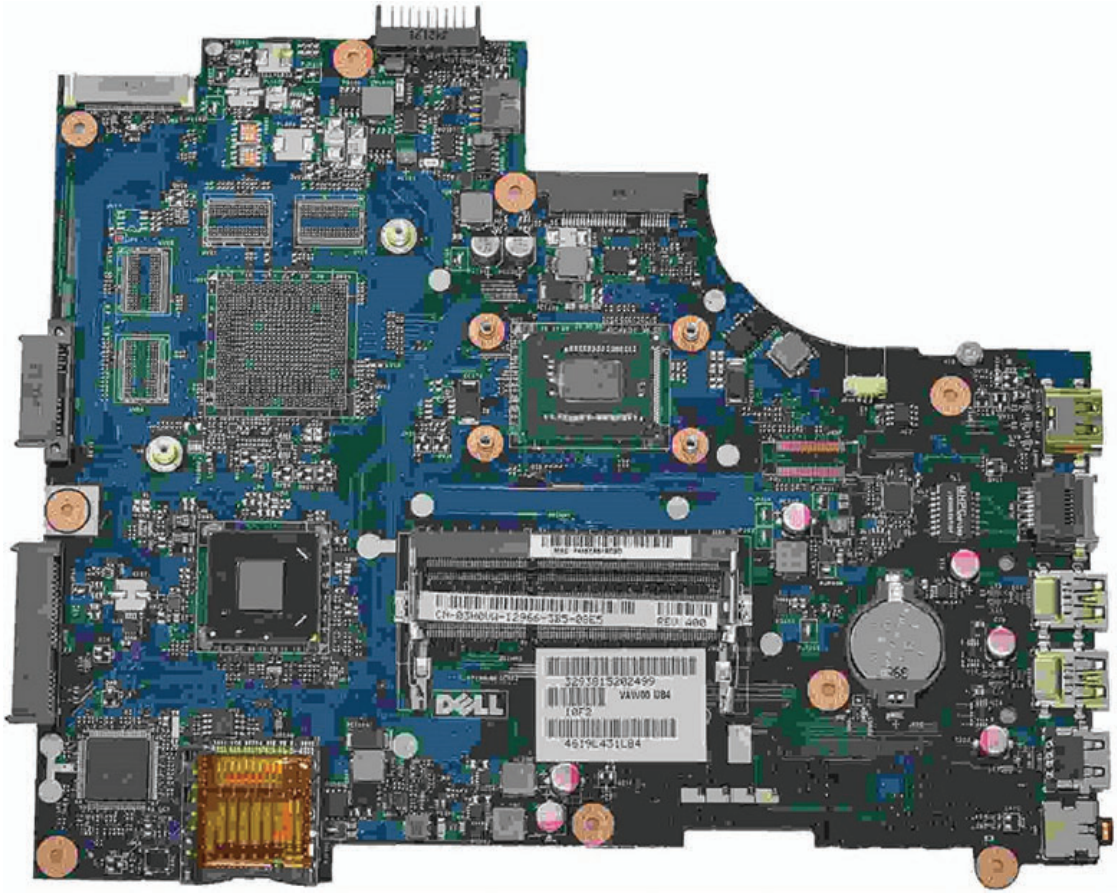

Nearly all laptop processors are soldered onto the motherboard, so you don't have to worry about CPU socket compatibility. If the CPU dies, you replace the entire motherboard. We will cover laptop components more extensively in Volume 1 Chapter 4, "Laptop and Mobile Device Hardware."

Power Connectors

In addition to these sockets and slots on the motherboard, a special connector (the 24-pin white block connector shown in Figure 1.12) allows the motherboard to be connected to the power supply to receive power. This connector is where the ATX power connector (mentioned in Volume 1 Chapter 2 in the section "Understanding Power Supplies") plugs in.

FIGURE 1.12 An ATX power connector on a motherboard

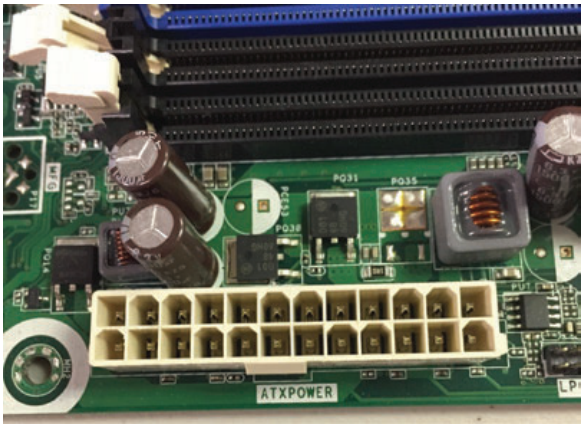

Onboard Nonvolatile Storage Connectors

Nearly all users store data, and the most widely used data storage device is a hard drive. Hard drives are great because they store data even when the device is powered off, which explains why they are sometimes referred to as nonvolatile storage. There are multiple types of hard drives, and we'll get into them in more detail in Volume 1 Chapter 2. Of course, those drives need to connect to the motherboard, and that's what we'll cover here.

Exam Objective 3.5 lists SATA, eSATA, SAN, and M.2 motherboard connectors. Know the differences between them!

Integrated Drive Electronics/Parallel ATA

At one time, *integrated drive electronics (IDE)* drives were the most common type of hard drive found in computers. Though often thought of in relation to hard drives, IDE was much more than a hard drive interface; it was also a popular interface for many other drive types, including optical drives and tape drives. Today, IDE is referred to as Parallel ATA (PATA), and it's a legacy technology that's obsolete. We mention it to help you understand where SATA (covered next, and a test objective) came from. And believe it or not, you might still see PATA connectors and drives in use today.

Figure 1.13 shows two PATA interfaces; you can see that one pin in the center is missing (as a key) to ensure that the cable gets attached properly. The industry now favors Serial ATA instead.

FIGURE 1.13 Two PATA hard drive connectors

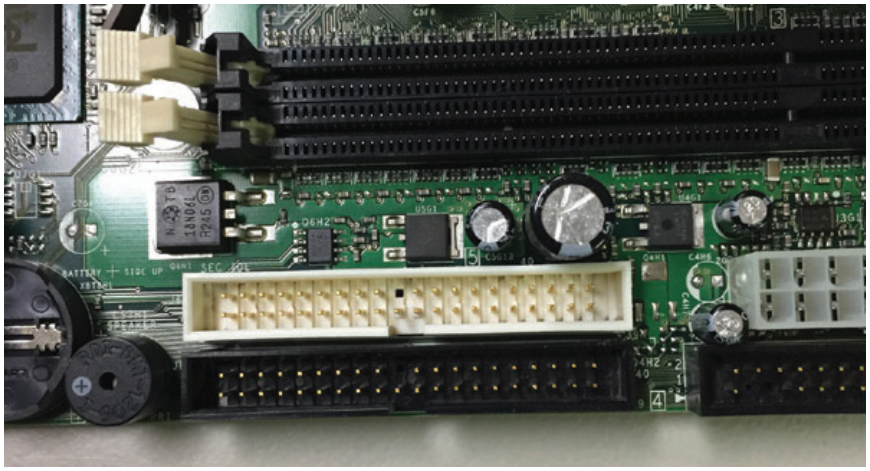

Serial ATA

Serial ATA (SATA) began as an enhancement to the original ATA specifications (IDE/PATA). Technology is proving that the orderly progression of data in a single-file path is superior to placing multiple bits of data in parallel and trying to synchronize their transmission to the point that all of the bits arrive simultaneously. Serial is simply faster than parallel in today's computer technology.

The first version of SATA, known as SATA 1.5 Gbps (and also by the less-preferred terms SATA I and SATA 150), used an 8b/10b-encoding scheme that requires 2 non-data overhead bits for every 8 data bits. The result is a loss of 20 percent of the rated bandwidth. The silver lining, however, is that the math becomes quite easy. Normally, you have to divide by 8 to convert bits to bytes. With 8b/10b encoding, you divide by 10. Therefore, the 150 Mbps throughput for which this version of SATA was nicknamed is easily derived as 1/10 of the 1.5 Gbps transfer rate. The original SATA specification also provided for hot swapping at the discretion of the motherboard and drive manufacturers.

Similar math works for SATA 3 Gbps, tagged as SATA II and SATA 300, and SATA 6 Gbps, which you might hear called SATA III or SATA 600. Note that each subsequent version doubles the throughput of the previous version. Figure 1.14 shows four SATA headers on a motherboard that will receive the data cable. Note that identifiers silkscreened onto motherboards often enumerate such headers. The resulting numbers are not related to the SATA version that the header supports. Instead, such numbers serve to differentiate headers from one another and to map to firmware identifiers, often visible within the BIOS configuration utility.

FIGURE 1.14 Four SATA headers

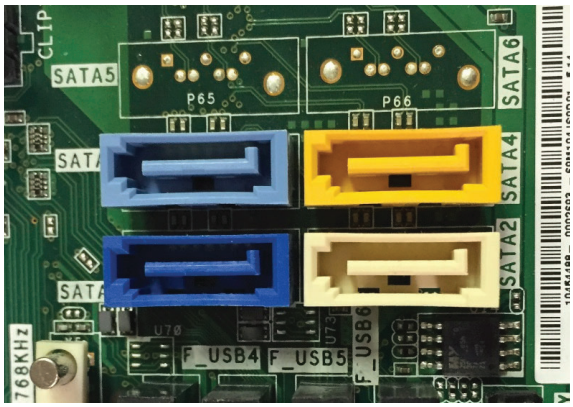

Another version of SATA that you will sometimes see is *external SATA (eSATA)*. As you might expect based on the name, this technology was developed for devices that reside outside of the case, not inside it. Some motherboards have an eSATA connector built in. If not, you can buy an expansion card that has eSATA ports and plugs into internal SATA connectors. Figure 1.15 shows an example of an eSATA port below two USB ports. SATA and eSATA standards are compatible. In other words, SATA 6 Gbps equals eSATA 6 Gbps.

FIGURE 1.15 eSATA port (bottom center)

SATA and eSATA ports (and PATA, for that matter) do not provide power, like USB and Thunderbolt do. Therefore, when connecting SATA and eSATA drives, a separate power connection is required in addition to the data cable. This limitation, along with slower speeds, has led to the demise of eSATA as a standard in favor of USB-C and Thunderbolt. We will cover both USB-C and Thunderbolt in more detail in Volume 1 Chapter 2.

What's in a Name?

The original *small computer system interface (SCSI)* standard was introduced in 1986 as a fast interface for hard drives and tape drives. Few motherboards had integrated SCSI controllers. If you wanted a SCSI hard drive, you needed to add an internal expansion card with a SCSI controller on it, which looked a bit like PATA but wider, and it had clips on either side. It was once popular for high-end systems, but the speed and lower cost of SATA made the original SCSI fade in popularity except for newer variants like iSCSI and *serial attached SCSI* (*SAS*; covered in Volume 1 Chapter 2), which are used in specialized applications.

As for the acronym pronunciation, the legend is that the developer wanted to call it *sexy*. After all, it was derived from the earlier *Shugart Associates System Interface (SASI)* standard which was called *sassy*. The *sexy* moniker was deemed too edgy for the computer industry at the time, and the *scuzzy* alternative took over. Who knows if the legend is true, but it makes a good story regardless.

M.2

The most recent development in expansion connections is *M.2* (pronounced "M dot 2"). So far it's primarily used for hard drives, but other types of devices, such as Wi-Fi, Bluetooth, Global Positioning System (GPS), and near-field communication (NFC) adapters are built for M.2 as well. We will cover M.2 in more depth in Volume 1 Chapter 2, thanks to how important it is to storage solutions.

It's important to call out that M.2 is a form factor, not a bus standard. The form factor supports existing SATA, USB, and PCIe buses. This means that if you hook up a SATA device to an M.2 slot (with the appropriate connector), the device speed will be regulated by SATA standards. Figure 1.16 shows an M.2 connector, highlighted and in the center pointing downward. To get to it, you will need to loosen a screw that holds a protective cover in place over the slot and remove the cover. The protective cover is shown in Figure 1.17, with the screw on the left.

M.2 is closely associated with hard drives, and the M.2 form factor is also in CompTIA A+ Exam Objective 3.4, "Compare and contrast storage devices." Because of that, we will go into more depth about M.2 and hard drive data transfer speeds in Volume 1 Chapter 2 when we discuss solid-state drives (SSDs).

Motherboard Headers

From the time of the very first personal computer, there has been a minimum expectation as to the buttons and LEDs that should be easily accessible to the user. At first, they generally appeared on the front of the case. In today's cases, buttons and LEDs have been added and placed on the top of the case or on a beveled edge between the top and the front. They have

also been left on the front or have been used in a combination of these locations. These buttons and lights, as well as other external connectors, plug into the motherboard through a series of pins known as *headers*. Examples of items that are connected using a header include the following:

- Power button
- Power light
- Reset button
- Audio jacks
- USB ports

FIGURE 1.16 M.2 connector on motherboard

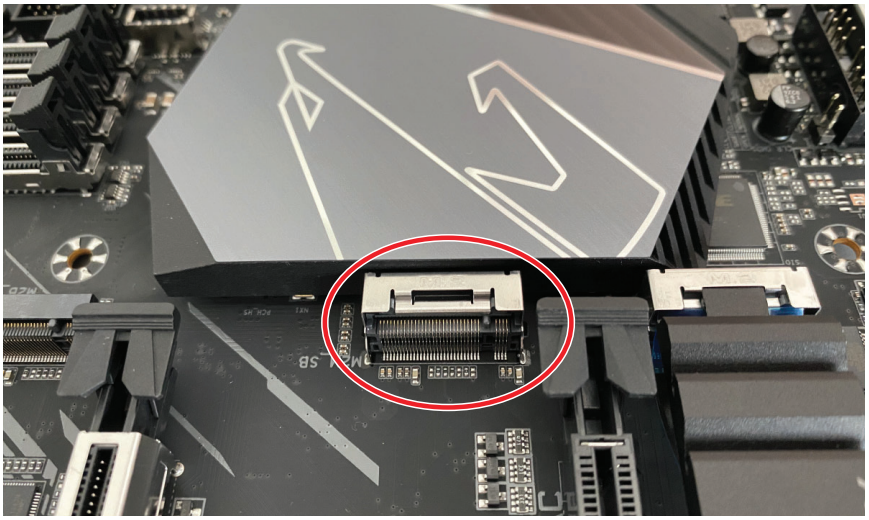

FIGURE 1.17 M.2 slot protective cover

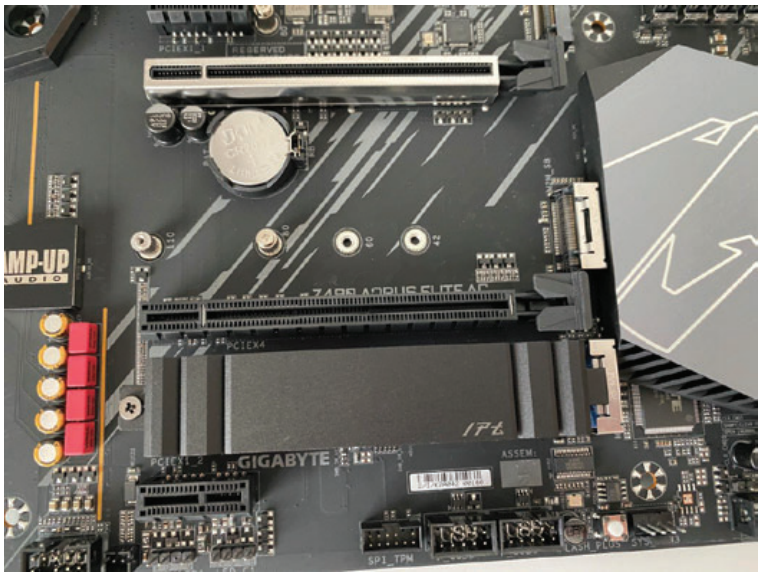

Headers for different connections are often spread throughout different locations on the motherboard—finding the right one can sometimes be a frustrating treasure hunt. Other headers are grouped together. For example, most of the headers for the items on the front or top panel of the case are often co-located. The purpose for the header will be printed on the motherboard, and while that may tell you what should connect there, it often lacks detail in how it should be connected. The motherboard manufacturer's website is a good place to go if you need a detailed diagram or instructions. Figure 1.18 shows several headers on a motherboard. On the left is a USB header, then a system fan header in the center, and a block of front-panel headers on the right, including the hard drive light, reset button, chassis intrusion detector, and power light.

FIGURE 1.18 Motherboard headers

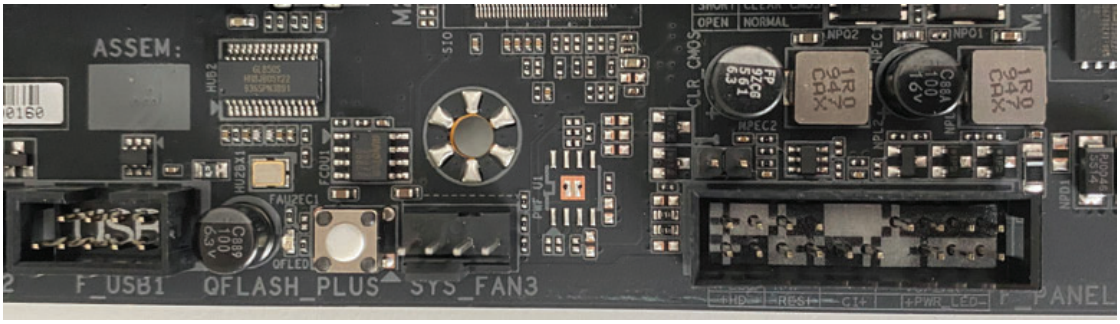

Exam Objective 3.5 says you need to know about motherboard headers. The exam objective doesn't list which headers you need to know, so we cover the most common ones here.

Power Button and Light

Users expect a *power button* to use to turn the computer on. (These were on the side or back of very early PCs.) The soft power feature available through the front power button, which is no more than a relay, allows access to multiple effects through the contact on the motherboard, based on how long the button is pressed. These effects can be changed through the BIOS/UEFI settings or operating system. Users also expect a *power light*, often a green LED, to assure them that the button did its job.

Reset Button

The *reset button* appeared as a way to reboot the computer from a cold startup point without removing power from the components. Keeping the machine powered tends to prolong the life of the electronics affected by power cycling. Pressing the reset button also gets around software lockups because the connection to the motherboard allows the system to restart from the hardware level. One disadvantage to power cycling is that certain circuits, such as memory chips, might need time to drain their charge for the reboot to be completely successful. This is why there is always a way to turn the computer off as well.

Audio Jacks

Early generations of optical drives had to have a special cable attached to the rear of the drive, which was then attached to the sound card if audio CDs were to be heard through the speakers attached to the sound card. Sound emanating from a CD-ROM running an application, such as a game, did not have to take the same route and could travel through the same path from the drive as general data. The first enhancement to this arrangement came in the form of a front 3.5mm jack on the drive's faceplate that was intended for headphones but could also have speakers connected to it. The audio that normally ran across the special cable was rerouted to the front jack when something was plugged into it.

Many of today's motherboards have 10-position pin headers designed to connect to standardized front-panel audio modules. Some of these modules have legacy AC'97 analog ports on them, whereas others have high-definition (HD) audio connections. Motherboards that accommodate both have a BIOS/UEFI setting that enables you to choose which header you want to activate, with the HD setting most often being the default.

USB Ports

So many temporarily attached devices feature USB connectivity, such as USB flash drives and cameras, that front-panel connectivity is a must. Finding your way to the back of the system unit for a brief connection is hardly worth the effort in some cases. For many years, motherboards have supplied one or more 10-position headers for internal connectivity of front-panel USB ports. Because this header size is popular for many connectors, only nine positions tend to have pins protruding, while the 10th position acts as a key, showing up in different spots for each connector type to discourage the connection of the wrong cable. Figure 1.19 shows USB headers on a motherboard. The labels "USB56" and "USB78" indicate that one block serves ports 5 and 6, while the other serves ports 7 and 8, all of which are arbitrary, based on the manufacturer's numbering convention. In each, the upper-left pin is "missing," which is the key.

FIGURE 1.19 Two motherboard USB headers

BIOS/UEFI

Firmware is the name given to any software that is encoded in hardware, usually a read-only memory (ROM) chip, and it can be run without extra instructions from the operating system. Most computers, large printers, routers, and devices with no operating system use firmware in some sense. The best example of firmware is a computer's *Basic Input/Output System (BIOS)*, which is burned into a chip. Also, some expansion cards, such as graphics adapters, use their own firmware utilities for setting up and configuring different functions.

The BIOS chip, also referred to as the ROM BIOS chip, is one of the most important chips on the motherboard. This special memory chip contains the BIOS system software that boots the system and allows the operating system to interact with certain hardware in the computer in lieu of requiring a more complex device driver to do so. The BIOS chip is easily identified; if you have a brand-name computer, this chip might have on it the name of the manufacturer and usually the word *BIOS*. For clones, the chip usually has a sticker or printing on it from one of the major BIOS manufacturers (AMI, Phoenix, Award, Winbond, and others).

PC technology and security requirements started to outclass what the BIOS could support, and in 2002 a successor to the BIOS was released, called *Unified Extensible Firmware Interface (UEFI)*. The extensible nature of the UEFI supports key features, such as these:

- Graphical interface
- Larger hard drives (BIOS was limited to 2.2 TB)
- Better security such as Secure Boot and TPM
- Faster boot times

Even though on PCs BIOS technology is outdated, you still hear the term used quite a bit. While the technologies are different, in practice, the terms BIOS and UEFI are essentially interchangeable.

Figure 1.20 gives you an idea of what a modern BIOS/UEFI chip might look like on a motherboard. Despite the 1998 copyright on the label, which refers only to the oldest code present on the chip, this particular chip can be found on motherboards produced as late as 2009. Notice also the Reset CMOS jumper at the lower left and its configuration silkscreen at the upper left. You might use this jumper to clear the CMOS memory, discussed shortly, when an unknown password, for example, is keeping you out of the BIOS/UEFI configuration utility. The jumper in the photo is in the Reset position, not the normal operating position. System bootup is typically not possible in this state.

FIGURE 1.20 A BIOS chip on a motherboard

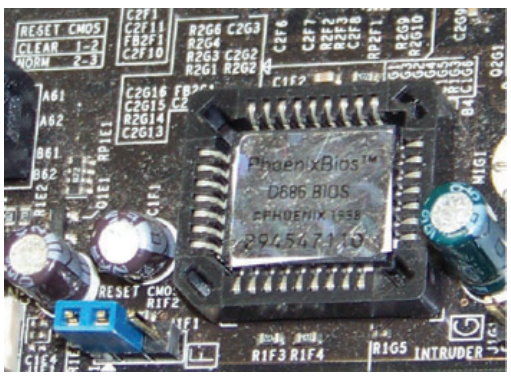

The BIOS/UEFI is responsible for several key functions related to the proper booting of a computer, and then it hands off control of the hardware to the operating system. Let's break some of these functions down into three groups: configuration and monitoring, security and encryption, and the power-on self-test (POST) routine.

Exam Objective 3.5 has quite a few BIOS/UEFI features to remember for the exam. They include boot options, USB permissions, TPM security features, fan considerations, Secure Boot, boot password, BIOS password, and temperature monitoring. Encryption is also part of this exam objective and covers TPM as well as HSM.

Configuration and Monitoring

The BIOS/UEFI contains basic configuration information about the CPU, RAM, and hard drives installed in a computer. It will detect the size and/or speed of these devices, and in some instances let you modify performance characteristics. Here are some features to be familiar with:

Boot Options In the BIOS/UEFI, you can set the device from which the computer boots an operating system. This can be a hard drive or an optical drive, or in some cases a flash drive. Figure 1.21 shows the boot option setting in a GIGABYTE AORUS motherboard.

FIGURE 1.21 UEFI boot option

USB Permissions USB ports can be physically disabled using the BIOS/UEFI interface. This can act as an additional security measure, because if the ports are disabled, someone can't introduce malware (such as a virus) into the computer by using a flash drive in the port.

Fan Considerations The BIOS/UEFI interface will allow you to turn system fans on or off as well as monitor the fan speeds. This is useful if you have a system that's overheating. It can also be helpful as a troubleshooting tool if you have a fan that's constantly speeding up or slowing down, or making excessive noise.

Temperature Monitoring The inside of a computer case can get hot, and the CPU is particularly good at generating excessive amounts of heat. While internal computer components are designed to withstand some heat, too much can cause them to malfunction. BIOS/UEFI utilities let you monitor temperature for the CPU, RAM, and often the ambient temperature of air inside the case as well.

Knowing how to enter the BIOS/UEFI and configure settings is an important skill for any technician. Exercise 1.1 shows you how to do it from Windows 11.

EXERCISE 1.1

Entering the UEFI from Windows 11

Some systems may let the user enter the BIOS/UEFI by pressing a key during the boot sequence. If this is the situation, a message will appear quickly on the screen, prior to the operating system loading screen, instructing the user to press a key such as F2, F12, or Delete.

Newer computers frequently omit this option. Computers running Windows 10 and 11 allow the user to reboot into the BIOS by using the Advanced startup option within the System Recovery tool.

1. In the Windows search bar, type **`Advanced startup`**. The Change advanced startup options system settings app will appear in the menu. Press Enter to open it. You will see a screen similar to the one in Figure 1.22.

FIGURE 1.22 System Recovery options

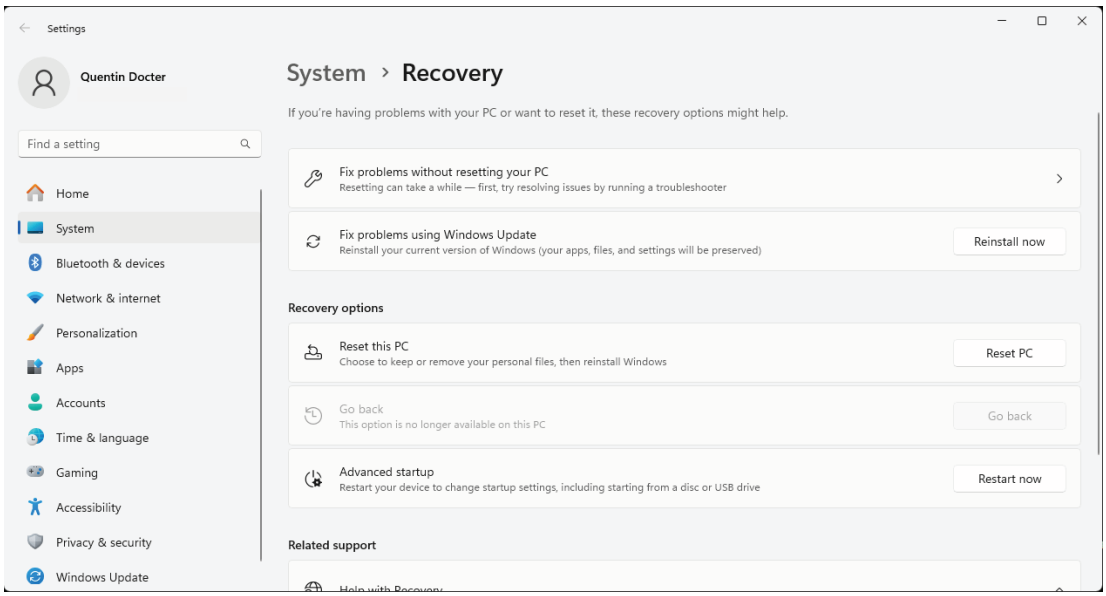
Settings
Quentin Docter
Find a setting
Home
System
Bluetooth & devices
Network & internet
Personalization
Apps
Accounts
Time & language
Gaming
Accessibility
Privacy & security
Windows Update
System › Recovery
If you're having problems with your PC or want to reset it, these recovery options might help.
Fix problems without resetting your PC
Resetting can take a while — first, try resolving issues by running a troubleshooter
Fix problems using Windows Update
Reinstall your current version of Windows (your apps, files, and settings will be preserved)
Reinstall now
Recovery options
Reset this PC
Choose to keep or remove your personal files, then reinstall Windows
Reset PC
Go back
This option is no longer available on this PC
Go back
Advanced startup
Restart your device to change startup settings, including starting from a disc or USB drive
Restart now
Related support

2. Next to Advanced Startup, click Restart Now.
3. After the system reboots, you will see the Choose An Option screen with three options: Continue To Windows 11, Troubleshoot, and Turn Off Your PC. Click Troubleshoot.
4. On the Troubleshoot screen, click Advanced Options.
5. On the Advanced Options screen (Figure 1.23), click UEFI Firmware Settings.

FIGURE 1.23 Advanced Options

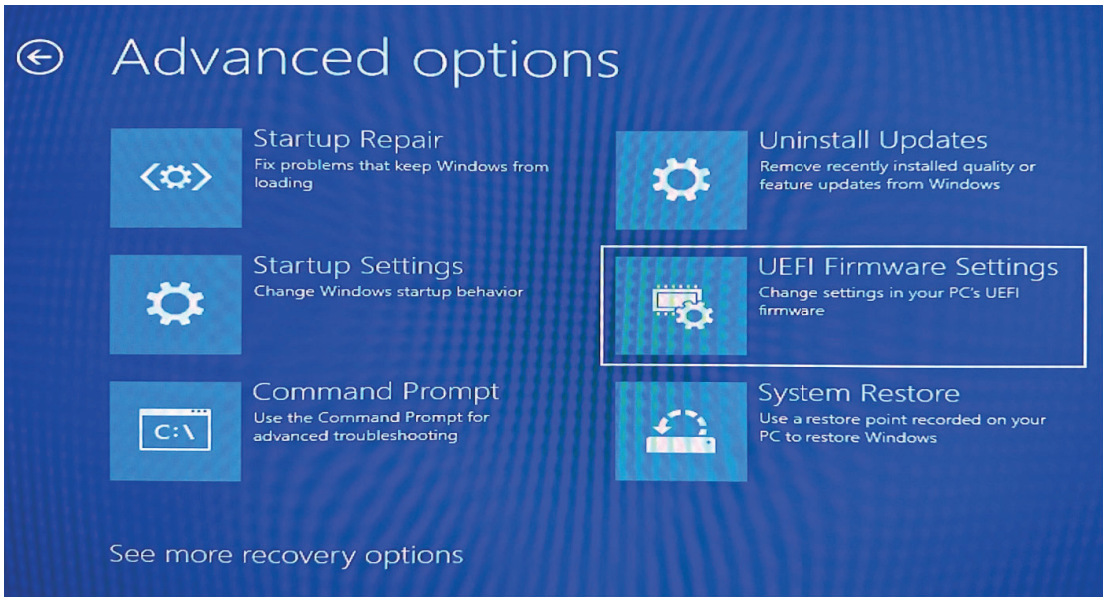
Advanced options
Startup Repair
Fix problems that keep Windows from loading
Uninstall Updates
Remove recently installed quality or feature updates from Windows
Startup Settings
Change Windows startup behavior
UEFI Firmware Settings
Change settings in your PC's UEFI firmware
Command Prompt
Use the Command Prompt for advanced troubleshooting
System Restore
Use a restore point recorded on your PC to restore Windows
See more recovery options

EXERCISE 1.1 *(continued)*

6. Click to confirm the reboot, and the system will reboot into the UEFI.

Explore the options; most UEFI interfaces have several tabs. Be careful, though, because changing settings in your UEFI could cause your computer to not operate properly. Don't change a setting unless you're sure you know what it does! If you do accidentally change something and you're concerned, there are two options. Either exit without saving, or reboot the computer using the power or reset button.

Security and Encryption

The BIOS/UEFI has always played a role in system security. Since the early days of the personal computer, the BIOS allowed the setting of two passwords—the user (or boot) password and the supervisor/administrator, or access, password. The *boot password* is required to leave the initial power-on screens and begin the process of booting an operating system. The administrator *BIOS password* is required before entering the BIOS/UEFI configuration utility. It is always a good idea to set the administrator password, but the boot password should not be set on public systems that need to boot on their own, in case of an unforeseen power cycle.

In more recent years, the role of the BIOS/UEFI in system security has grown substantially. BIOS/UEFI security has been extended to a point where the operating system is ready to take it over. The BIOS/UEFI was a perfect candidate to supervise security and integrity in a platform-independent way.

Most motherboards today have a dedicated security coprocessor, or cryptoprocessor, called the *Trusted Platform Module (TPM)*, installed on them. When enabled in the UEFI, UEFI works with the TPM and will boot only to an authenticated boot device. This authentication confirms that the hardware being booted to has been tied to the system containing the BIOS/UEFI and TPM, a process known as *sealing*. Sealing the devices to the system also prohibits the devices from being used after removing them from the system. For further security, the keys created can be combined with a PIN or password that unlocks their use or with a USB flash drive that must be inserted before booting.

In the AORUS UEFI, TPM is enabled in Settings ➢ Trusted Computing (Figure 1.24). The TPM settings are further shown in Figure 1.25.

FIGURE 1.24 UEFI settings

FIGURE 1.25 TPM settings

Microsoft's BitLocker uses the TPM to encrypt the entire drive. Normally, only user data can be encrypted, but BitLocker encrypts operating system files, the Registry, the hibernation file, and so on, in addition to those files and folders that file-level encryption secures. If any changes have occurred to the Windows installation, the TPM does not release the keys required to decrypt and boot to the secured volume. TPM status is shown in Windows 11 in the Device Security app. To open it, type **`tpm`** into the Windows search box and press Enter. Then click the Security Processor Details link and you will see a screen similar to the one shown in Figure 1.26.

Most motherboards come with a TPM chip installed, but if they don't, it's not possible to add one. In those situations, you can enable the same functionality by using a *hardware security module (HSM)*. An HSM is a security device that can manage, create, and securely store encryption keys—it enables users to safely encrypt and decrypt data. An HSM can take a few different forms. The simplest is a USB or PCIe device that plugs into a system. It could be set up for file encryption and decryption, required for the computer to boot, or both. For large-scale solutions, HSM-enabled servers can provide crypto services to an entire network. So while an HSM can perform the functionality of a TPM, it can have much broader application as well. They can be used together to form layers of security.

When a certain level of UEFI is used, the system firmware can also check digital signatures for each boot file it uses to confirm that it is the approved version and has not been tampered with. This technology is known as *Secure Boot*. An example of a BIOS/UEFI's

boot security screen is shown in Figure 1.27. The boot files checked include option ROMs (defined in the following section), the boot loader, and other operating system boot files. Only if the signatures are valid will the firmware load and execute the associated software.

FIGURE 1.26 Windows TPM configuration screen

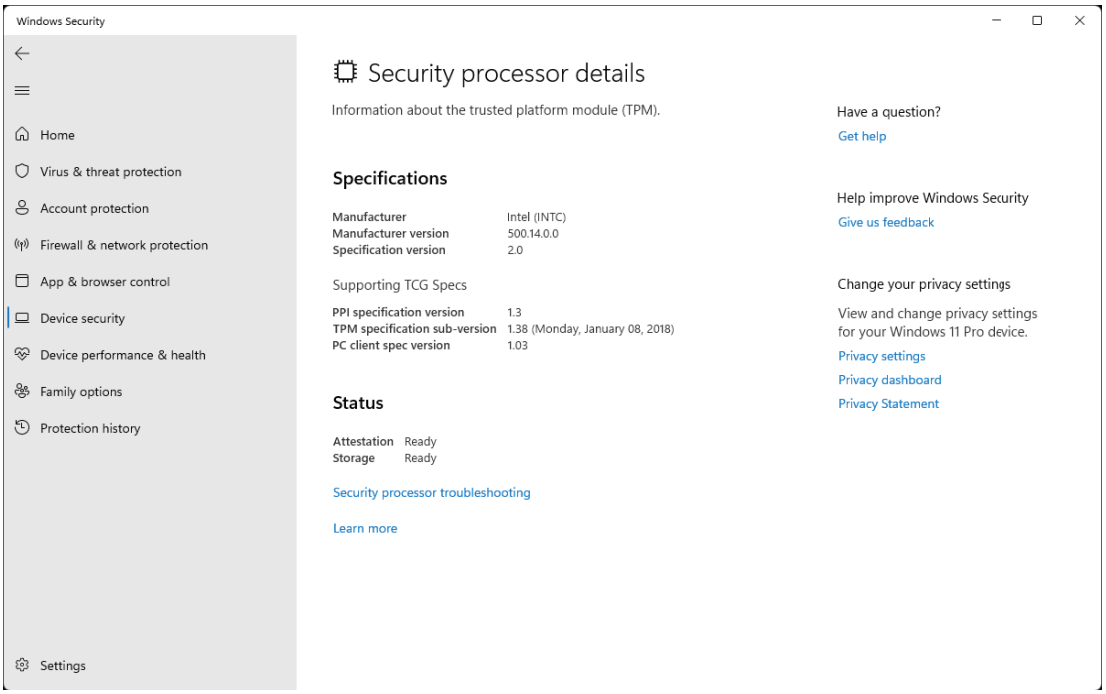

FIGURE 1.27 Secure Boot in UEFI

Some BIOS/UEFI firmware can monitor the status of a contact on the motherboard for intrusion detection. If the feature in the BIOS/UEFI is enabled and the sensor on the chassis is connected to the contact on the motherboard, the removal of the cover will be detected

and logged by the BIOS/UEFI. This can occur even if the system is off, thanks to the CMOS battery. At the next bootup, the BIOS/UEFI will notify you of the intrusion. No notification occurs over subsequent boots unless additional intrusion is detected.

POST Routine

A major function of the BIOS/UEFI is to perform a process known as a *power-on self-test (POST)*. POST is a series of system checks performed by the system BIOS/UEFI and other high-end components, such as the SCSI BIOS and the video BIOS, known collectively as *option ROMs*. Among other things, the POST routine verifies the integrity of the BIOS/UEFI itself. It also verifies and confirms the size of primary memory. During POST, the BIOS also analyzes and catalogs other forms of hardware, such as buses and boot devices, as well as manages the passing of control to the specialized BIOS/UEFI routines mentioned earlier. The BIOS/UEFI is responsible for offering the user a key sequence to enter the configuration routine as POST is beginning. Finally, once POST has completed successfully, the BIOS/UEFI selects the boot device highest in the configured boot order and executes the master boot record (MBR) or similar construct on that device so that the MBR can call its associated operating system's boot loader and continue booting up.

The POST process might end with a beep code or displayed code that indicates the issue discovered. Each BIOS/UEFI publisher has its own series of codes that can be generated. If you boot a computer and get a series of beeps and the operating system never loads, odds are there was a problem during the POST. At that point, check the motherboard manufacturer's documentation to see what the beep code is telling you.

Flashing the System BIOS/UEFI

If ever you find that a hardware upgrade to your system is not recognized, even after the latest and correct drivers (small programs that tell the OS how to interact with the hardware) have been installed, perhaps a BIOS/UEFI upgrade, also known as *flashing the BIOS*, is in order. Only certain hardware benefits from a BIOS/UEFI upgrade, such as drives and a change of CPU or RAM types. Very often, this hardware is recognized immediately by the BIOS/UEFI and has no associated driver that you must install. So, if your system doesn't recognize the new device, and there's no driver to install, the BIOS/UEFI is a logical target.

Let's be clear that we are not talking about entering the BIOS/UEFI setup utility and making changes to settings and subsequently saving your changes before exiting and rebooting. What we are referring to here is a replacement of the burned-in code within the BIOS itself. You might even notice after the upgrade that the BIOS setup utility looks different or has different pages and entries than before.

On older systems and certain newer ones, a loss of power during the upgrade results in catastrophe. The system becomes inoperable until you replace the BIOS/UEFI chip, if possible, or the motherboard itself. Most new systems, however, have a fail-safe or two. This could be a portion of the BIOS/UEFI that does not get flashed and has just enough code to

boot the system and access the upgrade image. It could be a passive section to which the upgrade is installed and switched to only if the upgrade is successful. Sometimes this is controlled onscreen. At other times, there may be a mechanism, such as a jumper, involved in the recovery of the BIOS/UEFI after a power event occurs. The safest bet is to make sure that your laptop has plenty of battery power and is connected to AC power or your desktop is connected to an uninterruptible power supply (UPS).

In all cases, if you think you need a BIOS/UEFI upgrade, do not consult the BIOS/UEFI manufacturer. Instead, go to the motherboard or system manufacturer and check its website. The motherboard or system manufacturer vendors have personalized their BIOS/UEFI code after licensing it from the BIOS/UEFI publisher. The vendor will give you access to the latest code as well as the appropriate flashing utility for its implementation.

CMOS and CMOS Battery

Your PC has to keep certain settings when it's turned off and its power cord is unplugged:

- Date
- Time
- Hard drive/optical drive configuration
- Memory
- CPU settings, such as overclocking
- Integrated ports (settings as well as enable/disable)
- Boot sequence
- Power management
- Virtualization support
- Security (passwords, Trusted Platform Module settings, etc.)

Your PC keeps these settings in a special memory chip called the *complementary metal oxide semiconductor (CMOS)* memory chip. Actually, CMOS (usually pronounced *see-moss*) is a manufacturing technology for integrated circuits. The first commonly used chip made from CMOS technology was a type of memory chip, the memory for the BIOS/UEFI. As a result, the term *CMOS* stuck and is the accepted name for this memory chip.

The BIOS/UEFI starts with its own default information and then reads information from the CMOS, such as which hard drive types are configured for this computer to use, which drive(s) it should search for boot sectors, and so on. Any overlapping information read from the CMOS overrides the default information from the BIOS/UEFI. A lack of corresponding information in the CMOS does not delete information that the BIOS knows natively. This process is a merge, not a write-over. CMOS memory is usually *not* upgradable in terms of its capacity and might be integrated into the BIOS/UEFI chip or the Southbridge.

Although there are technical differences, in the real world the terms BIOS/UEFI and CMOS (and BIOS/UEFI chip and CMOS chip) are used interchangeably. You will hear someone say, "Did you check the BIOS/UEFI?" or, "Did you check CMOS settings?" and they're referring to the same thing. Even though it's an oversimplification, most people feel comfortable with thinking about the BIOS/UEFI as the firmware that controls the boot process, whereas CMOS is the chip itself. It's incredibly unlikely that the A+ exam or any person is going to ask you to differentiate between the two, unless you get a job focusing specifically on those technologies.

To keep its settings, integrated circuit-based memory must have power constantly. When you shut off a computer, anything that is left in this type of memory is lost forever. The CMOS manufacturing technology produces chips with very low power requirements. As a result, today's electronic circuitry is more susceptible to damage from electrostatic discharge (ESD). Another ramification is that it doesn't take much of a power source to keep CMOS chips from losing their contents.

To prevent CMOS from losing its rather important information, motherboard manufacturers include a small battery called the *CMOS battery* to power the CMOS memory, shown to the left of the CPU socket in Figure 1.28. The batteries come in different shapes and sizes, but they all perform the same function. Most CMOS batteries look like large watch batteries or small cylindrical batteries. Today's CMOS batteries are most often of a long-life, nonrechargeable lithium chemistry.

FIGURE 1.28 CMOS battery

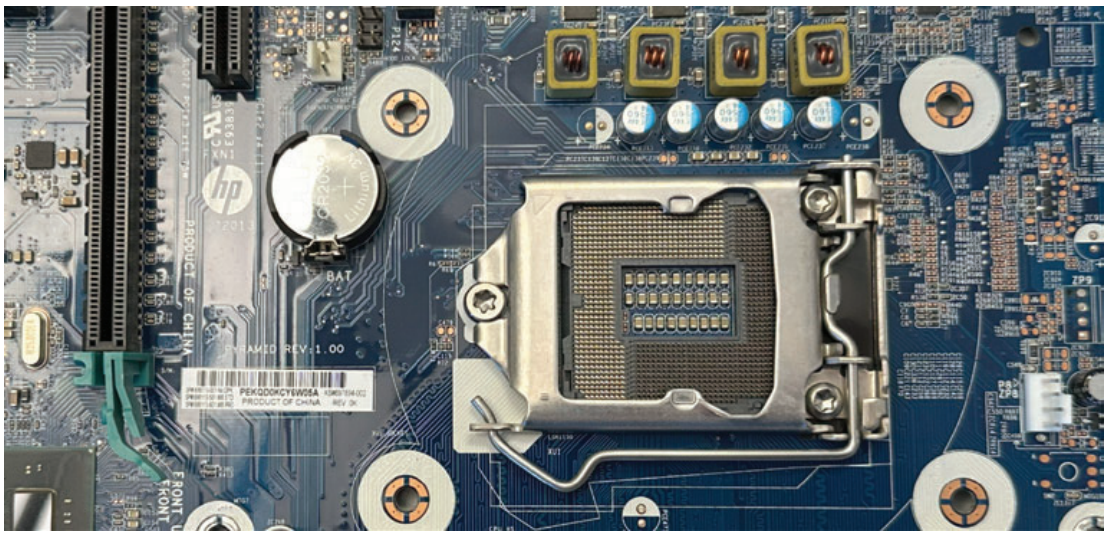

Understanding Processors

Now that you've taken a tour of the motherboard, you need to learn about the most important component on the motherboard: the CPU. The role of the CPU, or central processing unit, is to control and direct all the activities of the computer using both external and internal buses. From a technical perspective, the job of the CPU is to process, or do math, on large strings of binary numbers—0s and 1s. It is a processor chip consisting of an array of *billions* of transistors. The 13th-generation Intel Raptor Lake CPU has an estimated 12 billion.

Intel and Advanced Micro Devices, Inc. (AMD) are the two largest PC-compatible CPU manufacturers. Their chips were featured earlier in Table 1.1 during the discussion of the sockets into which they fit. From a compatibility standpoint, the most important thing to remember is that the motherboard and processor need to be made for each other. The rest of the hardware plugs into the motherboard and will be CPU brand agnostic.

The term *chip* has grown to describe the entire package that a technician might install in a socket. However, the word originally denoted the silicon wafer that is generally hidden within the carrier that you actually see. The external pins that you see are structures that can withstand insertion into a socket and are carefully threaded from the wafer's minuscule contacts. Just imagine how fragile the structures must be that you *don't* see.

Very early PC processors were manufactured in a rectangle with two rows of 20 pins known as a *dual in-line package (DIP)*—see Figure 1.29. There are still integrated circuits that use the DIP form factor; it's just no longer used for PC CPUs. Most modern CPUs use the LGA form factor although you will still see a fair number of PGA chips out there. Figure 1.9, earlier in this chapter, shows a PGA socket next to an LGA socket.

FIGURE 1.29 DIP and PGA

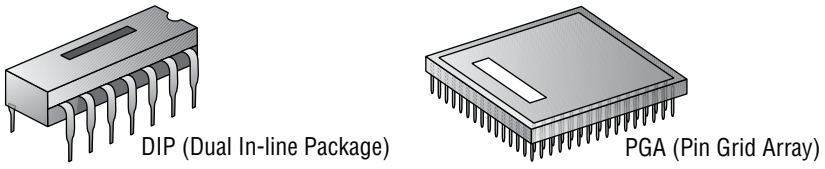

Intel and AMD both make extensive use of an inverted socket/processor combination of sorts. As mentioned earlier, the LGA packaging calls for the pins to be placed on the motherboard, while the mates for these pins are on the processor packaging. As with PGA, LGA is named for the landmarks on the processor, not the ones on the motherboard. As a result, the grid of metallic contact points, called *lands*, on the bottom of the CPU gives this format its name.

Figure 1.30 shows the top and bottom of an LGA CPU and a PGA CPU. Notice the lands on the left one and the pins on the right. Even though these CPUs are several years old now, the physical structure, outside of the number of pins, hasn't changed much. A decade-old CPU will look very much like a new one.

FIGURE 1.30 Top and bottom of LGA Intel and PGA AMD CPUs

Finding the CPU location in a computer is generally easy to do because it will be covered with a large heat sink, fan, or both. You likely won't be able to see the CPU itself because it's completely covered by the cooling system. The CPU is almost always located very close to the RAM to improve system speed, as shown in Figure 1.1, Figure 1.2, Figure 1.6, and Figure 1.10.

CPU Architecture

As noted in the previous section, the functional job of the processor is to do math on very large strings of 0s and 1s. How the CPU goes about doing that depends on its architecture. For commonly used processors, there are two major categories—those based on *Complex Instruction Set Computing (CISC)* and those based on *Reduced Instruction Set Computer (RISC)*.

Be familiar with the differences between x86/x64 and ARM CPUs as well as core configurations for the A+ exam, as part of Exam Objective 3.5.

x64/x86

CISC (pronounced like *disk*, but with a soft "c" like "cider") and RISC (pronounced *risk*) are examples of an instruction set architecture (ISA). Essentially, it's the set of commands that the processor can execute. Both types of chips, when combined with software, can ultimately perform all the same tasks. They just go about it differently. When programmers develop code, they develop it for a CISC or a RISC platform.

As the CISC name implies, instructions sent to the computer are relatively complex (as compared to RISC), and as such they can do multiple mathematical tasks with one instruction, and each instruction can take several clock cycles to complete. We'll talk more about CPU speeds in the "CPU Characteristics" section later, but for now, know that if a CPU is advertised as having 3.8 GHz speed, that means it can complete roughly 3.8 billion cycles in one second. The core of a processor can only do one thing at a time—it just does them very, very quickly so it looks like it's multitasking.

CISC was the original ISA for microprocessors, and the most well-known example of CISC technology is the x64/x86 platform popularized by Intel. AMD processors are CISC chips as well. So where did the terms x64 and x86 come from? First, just a bit more theory.

There is a set of data lines between the CPU and the primary memory of the system—remember the *bus*? The most common bus today is 64 bits wide, although there are still some 32-bit buses kicking around out there. In older days, buses could theoretically be as narrow as 2 bits, and 8-bit and 16-bit buses were popular for CPUs for several years. The wider the bus, the more data that can be processed per unit of time, and hence, more work can be performed. Internal registers in the CPU might be only 32 bits wide, but with a 64-bit system bus, two separate pipelines can receive information simultaneously. For true 64-bit CPUs, which have 64-bit internal registers and can run x64 versions of operating systems, the external system data bus will always be 64 bits wide or some larger multiple thereof.

In the last paragraph we snuck in the term *x64*, and by doing that we also defined it. It refers to processors that are designed to work with 64 bits of data at a time. To go along with it, the operating system must also be designed to work with x64 chips.

Contrast that with processors that can handle only 32 bits of information at once. Those are referred to as *x86* processors. You might look at that last sentence and be certain that we made a typo, but we didn't. For a long time when 32-bit processors were the fastest on the PC market, Intel was the dominant player. Their CPUs had names like 80386 (aka i386) and 80486 (i486) and were based on the older 16-bit 80286 and 8086. Since the i386 and i486 were the most popular standards, the term x86 sprung up to mean a 32-bit architecture. So even though it may seem counterintuitive due to the numbers, x64 is newer and faster than x86.

Remember that x64/x86 technology is typically found in PCs, laptops, and servers. As you'd expect, 64-bit CPUs (x64) can handle much larger amounts of data and RAM than 32-bit CPUs (x86) can.

Advanced RISC Machine

Moving into the RISC architecture, the primary type of processor used today is known as an *Advanced RISC Machine (ARM)* CPU. Depending on who you talk to and which sources you prefer, there are conflicting stories about if that's actually the right acronym, as ARM is also known as an Acorn RISC Machine. Regardless of what it stands for, ARM is a competing technology to Intel and AMD x64-based CPUs.

Based on the RISC acronym, one might think that the reduced set of instructions the processor can perform makes it inferior somehow, but that's not the case. Tasks just need to get executed in different ways. To use a human example, let's say that we tell you to add the number 7 to itself seven times. One way to do that is to use one step of multiplication: 7 × 7 equals 49. But what if we said you can't use multiplication? You can still get to the answer by using addition. It will just take you seven steps instead of one. Same answer, different process. That's kind of how RISC compares to CISC.

RISC processors have some advantages over their CISC counterparts. They can be made smaller than CISC chips and they produce less heat, making them ideal for mobile devices. In fact, nearly all smartphones use RISC-based chips, such as Apple's A18 and Samsung's Snapdragon and Exynos series processors. On the downside, RISC processors use more memory than CISC ones do because it takes more code to complete a task with a RISC chip. Figure 1.31 shows an older Samsung RISC processor used in a smartphone. It's about the same size as a PC CPU. Also notice that the connector isn't standard like LGA or PGA; rather, it's developed specifically for the device using it.

FIGURE 1.31 Top (a) and bottom (b) of a Samsung RISC CPU

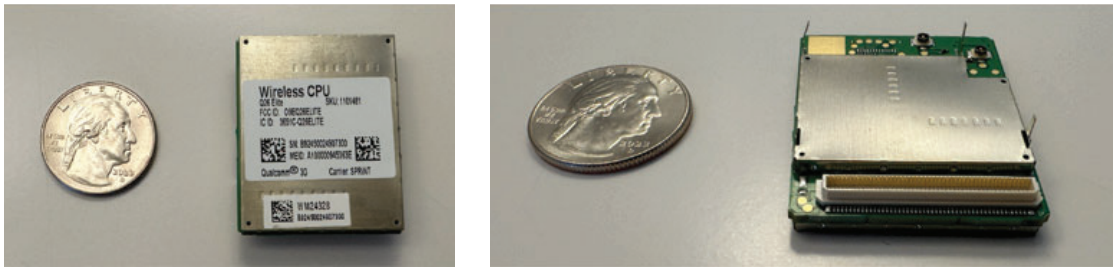

(a) (b)

Like x64/x86, ARM processors have evolved over time—64-bit implementations are the most current, and they are designated ARM64; 32-bit versions are known simply as ARM.

Remember that ARM chips are found mostly in smartphones, tablets, and smaller mobile devices. RISC chips handle smaller, highly optimized sets of instructions when compared to x64/x86 processors.

CPU Cores

Older processors were *single-core*, meaning that there was one set of instruction pathways through the processor. Therefore, they could process one set of tasks at a time. Designers then figured out how to speed up the computer by creating multiple cores within one processor package. Each core effectively operates as its own independent processor, provided that the operating system and applications are able to support *multicore* technology. (Nearly all do.)

Today, almost all desktop CPUs in the market are multicore. The number of cores you want may determine the processor to get. For example, the 13th-generation Intel Core i7 has 16 cores whereas the i5 has 14. Newer Intel processors will also designate two types of cores, P-cores, or performance cores, and E-cores, or efficient cores. P-cores are faster and support *hyperthreading* (the ability to manage multiple threads of an application at once), whereas E-cores are smaller and more heat-efficient. Ryzen processors will differ in core counts as well—the Ryzen 9 9950X has 16 cores whereas the Ryzen 7 9700X has 8 cores. AMD does not currently designate core types like Intel does.

As an interesting aside, throughout the 1980s and 1990s, processor clock speeds grew at very fast rates. But they haven't gotten much faster over the last 15–20 years as designers ran into technical limitations for increasing speeds. Instead, computers have gotten faster because of the introduction of multicore processors, which is basically like having four (or eight or 16) separate processors in one computer.

CPU Characteristics

When looking for a processor, you might have several decisions to make. Do you want an Intel or an AMD CPU? Which model? How fast should it be? What features does it need to support? In this section, we will take a look at some characteristics of processor performance.

Speed

The speed of the processor is generally described in clock frequency. All PC CPUs today are rated in gigahertz (GHz), or billions of cycles per second, and older ones in megahertz (MHz), or millions of cycles per second. Since the dawn of the personal computer industry, motherboards have included oscillators, which are quartz crystals shaved down to a specific geometry so that engineers know exactly how they will react when a current is run through them. The phenomenon of a quartz crystal vibrating when exposed to a current is known as the *piezoelectric effect*. The crystal (XTL) known as the *system clock* keeps the time for the flow of data on the motherboard. How the front-side bus uses the clock leads to an *effective*

clock rate known as the *FSB speed*. As discussed in the section "Types of Memory," later in this chapter, the FSB (front-side bus) speed is computed differently for different types of RAM (DDR5, DDR4, and so forth). From here, the CPU multiplies the FSB speed to produce its own internal clock rate, producing the *CPU speed* mentioned at the beginning of this paragraph.

The speed of a processor can also be tweaked by *overclocking*, or running the processor at a higher speed than the one at which the manufacturer rated it. Running at a higher speed requires more voltage and also generates more heat, which can shorten the life of the CPU. Manufacturers have historically discouraged the practice (of course, they want you to just buy a faster and more expensive CPU) because it can cause a CPU to have a shorter lifespan, and it may void any warranty. However, CPUs are sold today that specifically give you the ability to overclock them. Our official recommendation is to not do it unless the manufacturer says it's okay. If you're curious, plenty of information on how to overclock is available online.

Multithreading and Hyperthreading

The string of instructions that a CPU runs is known as a *thread*. Old processors were capable of running only one thread at a time, whereas newer ones can run multiple threads at once. This is called *multithreading*.

Intel markets their multithreading technology as *Hyper-Threading Technology (HTT)*. HTT is a form of simultaneous multithreading (SMT). SMT takes advantage of a modern CPU's superscalar architecture. Superscalar processors can have multiple instructions operating on separate data in parallel.

HTT-capable processors appear to the operating system to be two processors. As a result, the operating system can schedule two threads at the same time, as in the case of symmetric multiprocessing (SMP), where two or more processors use the same system resources. In fact, the operating system must support SMP in order to take advantage of HTT. If the current thread stalls because of missing data caused by, say, cache or branch prediction issues, the execution resources of the processor can be reallocated for a different thread that is ready to go, reducing processor downtime.

HTT manifests itself in the Windows 11 Task Manager by, for example, showing the number of logical processors as being twice as many as the number of cores the CPU has (see Figure 1.32).

For an in-market example, compare the 13th Generation Intel Core CPUs in Table 1.3. Remember that Intel E-cores (efficient cores) do not support multithreading, whereas the P-cores (performance cores) can run two threads each. So not only does the i7 have two more cores than the i5, those two cores are P-cores, so the i7 can run four more threads at one time. The i9 is even one further step up from the i7. For everyday email and Internet use, the differences won't amount to much. But for someone who is using resource-intensive apps such as online gaming or virtual reality, the differences, especially between i5 and i9 processors, can be important.

FIGURE 1.32 CPU performance in Task Manager

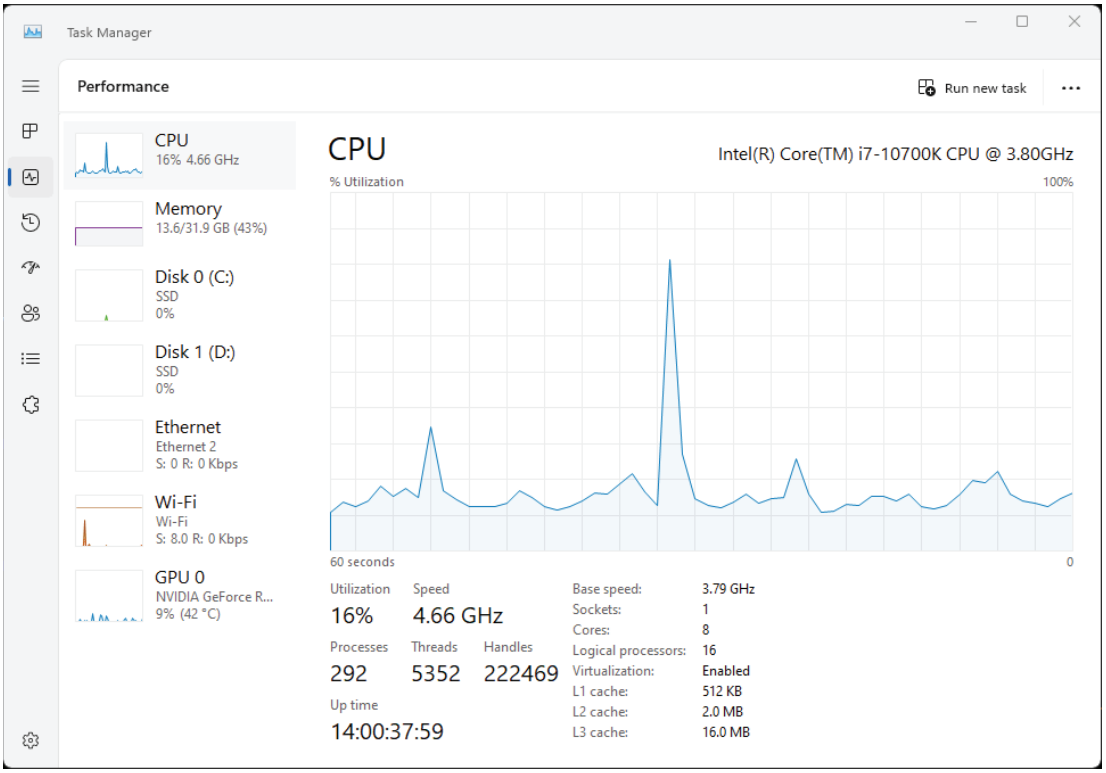

TABLE 1.3 13th Generation Intel Core CPU comparison

CPU number	P-cores	E-cores	Max CPU threads
i5-13600KF	6	8	20
i7-13700KF	8	8	24
i9-13900KF	8	16	32

Real World Scenario

Which CPU Do You Have?

There are a few easy ways to find out which processor a computer has, without cracking the case and ripping the heat sink off the CPU itself. One is to look in Control Panel ➢ System to get the About screen, as shown in Figure 1.33. Even more detailed information can be found by running the System Information app by typing **msinfo** into the Windows search box and

pressing Enter. Of course, third-party utilities such as CPU-Z that we showed you earlier will work too. A final way to determine a computer's CPU is to save your work, exit any open programs, and restart the computer. Watch closely as the computer boots back up. You should see a notation during the POST routine that tells you what chip the system is using.

FIGURE 1.33 System information in the About screen

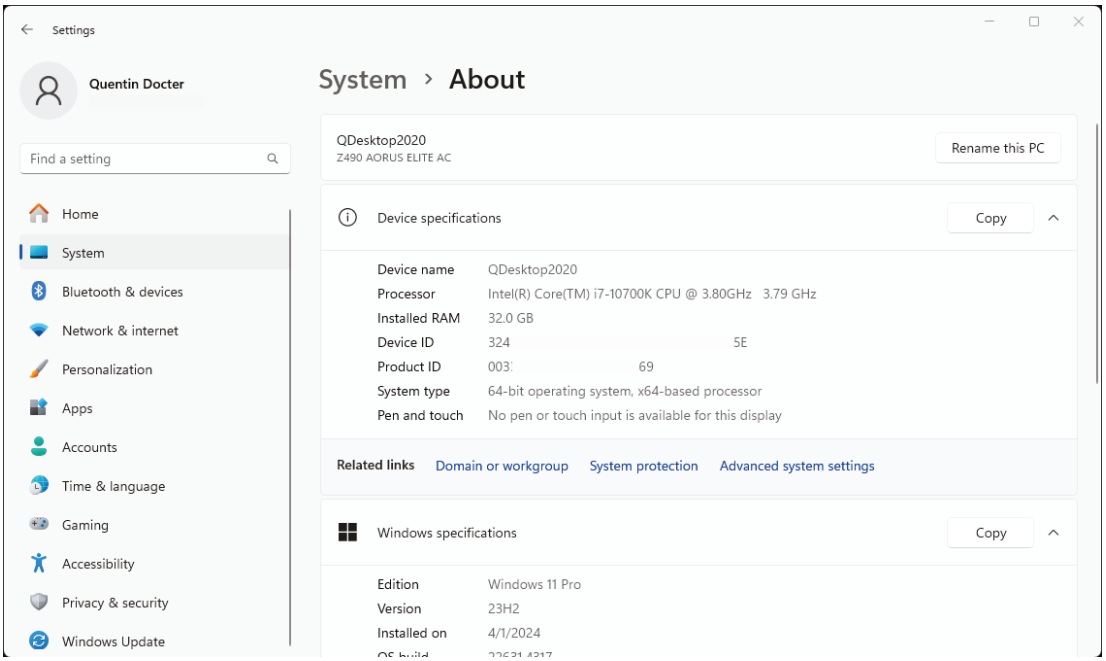

Virtualization Support

Hardware-based virtualization, which must be supported by the CPU, eases the burden on the system that software-based virtualization imposes. For more information on virtualization, see Volume 1 Chapter 11, "Virtualization and Cloud Computing." AMD calls its virtualization technology *AMD-V* (V for virtualization), whereas Intel calls theirs *Virtualization Technology (VT)*. Most processors made today support virtual technology, but not all. Keep in mind that the BIOS/UEFI and operating system must support it as well for virtualization to work. You may need to manually enable the virtualization support in the BIOS/UEFI before it can be used. If you have an Intel processor and would like to check its support of VT, visit the following site to download the Intel Processor Identification Utility:

```
www.intel.com/content/www/us/en/download/12136/intel-
processor-identification-utility-windows-version.html
```

Understanding Memory

"More memory, more memory, I don't have enough memory!" Adding memory is one of the most popular, easy, and inexpensive ways to upgrade a computer. As the computer's CPU works, it stores data and instructions in the computer's memory. Contrary to what you might expect from an inexpensive solution, memory upgrades tend to afford the greatest performance increase as well, up to a point. Motherboards have memory limits; operating systems have memory limits; CPUs have memory limits.

To identify memory visually within a computer, look for several thin rows of small circuit boards sitting vertically, potentially packed tightly together near the processor. In situations where only one memory stick is installed, it will be that stick and a few empty slots that are tightly packed together. Figure 1.34 shows where memory is located in a system—in this case, all four banks are full.

FIGURE 1.34 Four memory sticks in a motherboard

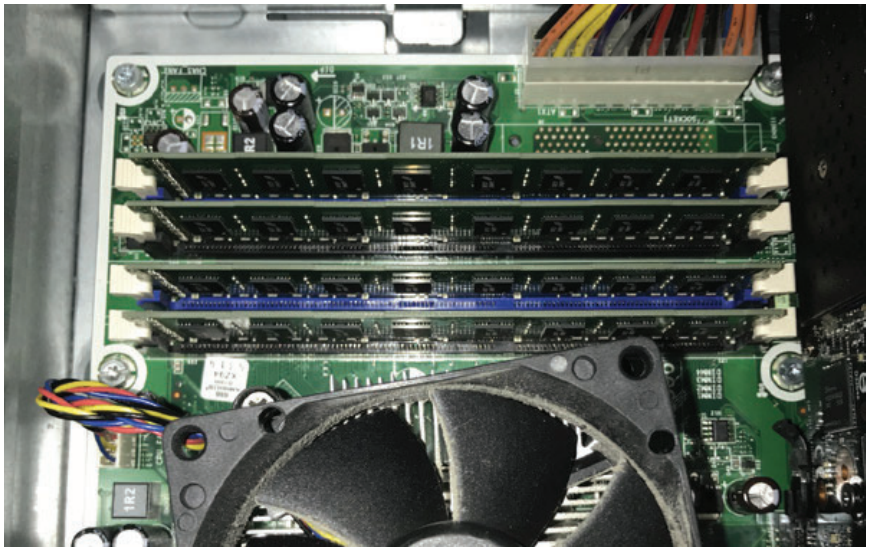

Important Memory Terms

There are a few technical terms and phrases that you need to understand with regard to memory and its function:

- Parity checking
- Error-correcting code (ECC)
- Channel configurations

These terms are discussed in detail in the following sections.

ECC will sometimes be referred to as error correction code as well. As with other acronyms in computing that might have multiple ways to spell it out, focus less on what it stands for and more on what it does.

A+ Exam Objective 3.3, "Compare and contrast RAM characteristics," calls out error-correcting code (ECC) and channel configurations as topics to know.

Parity Checking

Parity checking is a rudimentary error-checking scheme offered in older RAM that offers no error correction. Parity checking works most often on a byte, or 8 bits, of data. A ninth bit is added at the transmitting end and removed at the receiving end so that it does not affect the actual data transmitted. If the receiving end does not agree with the parity that is set in a particular byte, a parity error results, and the data is deemed unusable. Systems that do not support or reserve the location required for the parity bit are said to implement *non-parity memory*.

In the early days of personal computing, almost all memory was parity-based. As quality has increased over the years, parity checking in the RAM subsystem has become more uncommon.

Error Checking and Correcting

The next step in the evolution of memory error detection is known as *error-correcting code (ECC)*. If memory supports ECC, check bits are generated and stored with the data. An algorithm is performed on the data and its check bits whenever the memory is accessed. If the result of the algorithm is all zeros, then the data is deemed valid and processing continues. ECC can detect single- and double-bit errors and actually correct single-bit errors. In other words, if a particular byte—group of 8 bits—contains errors in 2 of the 8 bits, ECC can recognize the error. If only 1 of the 8 bits is in error, ECC can correct the error.

ECC RAM is essential for environments where critical or highly sensitive data is used. Consequently, you will see ECC RAM used a lot in servers, but it's uncommon in end-user PCs. ECC RAM is also more expensive than non-ECC RAM. But as more than one server administrator has said, "If it's important enough to be on a storage server, it's important enough to use ECC RAM."

For ECC RAM to perform error checking and correcting, the motherboard it's installed on must support ECC. If ECC RAM is installed in a motherboard that does *not* support ECC, the RAM should still be recognized and function—it just won't have error checking. The opposite is not true, however. Adding non-ECC RAM to an ECC motherboard can damage the ECC functionality of the motherboard.

Channel Configurations

Standard memory controllers manage access to memory in chunks of the same size as the system bus's data width. This is considered communicating over a single channel. Most modern processors have a 64-bit system data bus. This means that a standard memory controller can transfer exactly 64 bits of information at a time. Communicating over a single channel is a bottleneck in an environment where the CPU and memory can both operate faster than the conduit between them. Up to a point, every channel added in parallel between the CPU and RAM serves to ease this constriction.

Memory controllers that support dual-channel and greater memory implementation were developed in an effort to alleviate the bottleneck between the CPU and RAM. *Dual-channel memory* is the memory controller's coordination of two memory banks to work as a synchronized set during communication with the CPU, doubling the specified system bus width from the memory's perspective. *Triple-channel memory*, then, demands the coordination of three memory modules at a time. *Quad-channel memory* is the coordination of four memory modules at once. Collectively, they are known as *multichannel memory* implementations.

Because today's processors largely have 64-bit external data buses, and because one stick of memory satisfies this bus width, there is a 1:1 ratio between banks and modules. This means that implementing multichannel memory in today's most popular computer systems requires that two, three, or four memory modules be installed at a time. Note, however, that it's the motherboard, not the memory, that implements multichannel memory (more on this in a moment). *Single-channel memory*, in contrast, is the classic memory model that dictates only that a complete bank be satisfied whenever memory is initially installed or added. One bank supplies only half the width of the effective bus created by dual-channel support, for instance, which by definition pairs two banks at a time.

In almost all cases, multichannel implementations support single-channel installation, but poorer performance should be expected. Multichannel motherboards often include slots of different colors, usually one of each color per set of slots. To use only a single channel, you populate slots of the same color, skipping neighboring slots to do so. Filling neighboring slots in a dual-channel motherboard takes advantage of its dual-channel capability.

Because of the special tricks that are played with memory subsystems to improve overall system performance, care must be taken during the installation of disparate memory modules. In the worst case, the computer will cease to function when modules of different speeds or different capacities are placed together in slots of the same channel. If all of these parameters are identical, there should be no problem with pairing modules. Nevertheless, problems could still occur when modules from two different manufacturers or certain unsupported manufacturers are installed, all other parameters being the same. Technical support or documentation from the manufacturer of your motherboard should be able to help with such issues.

Although it's not the make-up of the memory that leads to multichannel support but instead the technology on which the motherboard is based, some memory manufacturers still package and sell pairs and triplets of memory modules in an effort to give you peace

of mind when you're buying memory for a system that implements multichannel memory architecture. Keep in mind, the motherboard memory slots have the distinctive color-coding, not the memory modules.

I Can't Fill All My Memory Slots

All motherboard manufacturers document the quantity and types of modules that their equipment supports. Consult your documentation when you have questions about supported memory. Most manufacturers require that slower memory be inserted in lower-numbered memory slots. This is because such a system adapts to the first module it sees, looking at the lower-numbered slots first. Counterintuitively, however, it might be required that you install modules of larger capacity rather than smaller modules in lower-numbered slots.

Additionally, memory technology continues to advance after each generation of motherboard chipsets is announced. Don't be surprised when you attempt to install a single module of the highest available capacity in your motherboard and the system doesn't recognize the module, either by itself or with others. That capacity of module might not have been in existence when the motherboard's chipset was released. Sometimes, flashing the BIOS is all that is required. Other times it just won't work. Consult the motherboard's documentation.

Finally, know that dual- and quad-channel memory support is common today, but triple-channel is less so. As with anything else, check the motherboard's documentation to see what it supports.

Types of Memory

Memory comes in many formats. Each one has a particular set of features and characteristics, making it best suited for a particular application. Some decisions about the application of the memory type are based on suitability; others are based on affordability to consumers or marketability to computer manufacturers. The following list gives you an idea of the vast array of memory types and subtypes—and this is only a partial list of what's been in the market:

- DRAM (dynamic random access memory)

  ADRAM (asynchronous DRAM)

  FPM DRAM (fast page mode DRAM)

  EDO DRAM (extended data out DRAM)

  BEDO DRAM (burst EDO DRAM)

SDRAM (synchronous DRAM)

SDR SDRAM (single data rate SDRAM)

DDR SDRAM (double data rate SDRAM)

DDR2 SDRAM (double data rate, version 2, SDRAM)

DDR3 SDRAM (double data rate, version 3, SDRAM)

DDR4 SDRAM (double data rate, version 4, SDRAM)

DDR5 SDRAM (double data rate, version 5, SDRAM)

- SRAM (static random access memory)
- ROM (read-only memory)

Pay particular attention to all synchronous DRAM types (such as DDR4 and DDR5) as that's the most common type in use. Note that the type of memory does not dictate the packaging of the memory. Conversely, however, you might notice one particular memory packaging holding the same type of memory every time you come across it. Nevertheless, there is no requirement to this end. Let's dive into details on some of these key memory types.

DRAM

DRAM is dynamic random access memory. This is what most people are talking about when they mention RAM. When you expand the memory in a computer, you are adding DRAM chips. You use DRAM to expand the memory in the computer because it's a cheaper type of memory. Dynamic RAM chips are cheaper to manufacture than most other types because they are less complex. *Dynamic* refers to the memory chips' need for a constant update signal (also called a *refresh signal*) in order to keep the information that is written there. If this signal is not received every so often, the information will bleed off and cease to exist. Currently, the most popular implementations of DRAM are based on synchronous DRAM and include DDR4 and DDR5. Occasionally you will see some older versions of DDR in the field, and new versions periodically arrive. For example, DDR6 may show up in 2025. It will take a while before it's widely adopted, however.

Asynchronous DRAM earned its name because its speed was not tied, or synchronized, with the motherboard's FSB. This required advanced BIOS configurations to make the memory play nice with the CPU. Implementations of ADRAM included FPM, EDO, and BEDO. None of these technologies are in use today.

Synchronous DRAM (SDRAM) is named as such because it shares a common clock signal with the computer's system-bus clock, which provides the common signal that all local-bus components use for each step that they perform. This characteristic ties SDRAM to the speed of the FSB and hence the processor, eliminating the need to configure the CPU to wait for the memory to catch up.

Originally, *SDRAM* was the term used to refer to the only form of synchronous DRAM on the market. As the technology progressed, and more was being done with each clock signal on the FSB, various forms of SDRAM were developed. What was once called simply SDRAM needed a new name retroactively. Today, we use the term *single data rate SDRAM (SDR SDRAM)* to refer to this original type of SDRAM.

SDR SDRAM

SDR SDRAM is a legacy RAM technology, and it is presented here only to provide a basis for the upcoming discussion of DDR and other more advanced RAM. With SDR SDRAM, every time the system clock ticks, 1 bit of data can be transmitted per data pin, limiting the bit rate per pin of SDRAM to the corresponding numerical value of the clock's frequency. With today's processors interfacing with memory using a parallel data-bus width of 8 bytes (hence the term *64-bit processor*), a 100 MHz clock signal produces 800 MBps. That's mega*bytes* per second, not mega*bits*. Such memory modules are referred to as PC100, named for the true FSB clock rate on which they rely. PC100 was preceded by PC66 and succeeded by PC133, which used a 133 MHz clock to produce 1,066 MBps of throughput.

Note that throughput in megabytes per second is easily computed as eight times the rating in the name. This trick works for the more advanced forms of SDRAM as well. The common thread is the 8-byte system data bus. Incidentally, you can double throughput results when implementing dual-channel memory.

DDR SDRAM

Double data rate (DDR) SDRAM earns its name by doubling the transfer rate of ordinary SDRAM; it does so by *double-pumping* the data, which means transferring a bit per pin on both the rising and falling edges of the clock signal. This obtains twice the transfer rate at the same FSB clock frequency. It's the increasing clock frequency that generates heating issues with newer components, so keeping the clock the same is an advantage. The same 100 MHz clock gives a DDR SDRAM system the impression of a 200 MHz clock compared to an SDR SDRAM system. For marketing purposes, and to aid in the comparison of disparate products (DDR vs. SDR, for example), the industry has settled on the practice of using this effective clock rate as the speed of the FSB.

Module Throughput Related to FSB Speed

There is always an 8:1 module-to-chip (or module-to-FSB speed) numbering ratio because of the 8 bytes that are transferred at a time with 64-bit processors (*not* because of the ratio of 8 bits per byte). The formula in Figure 1.35 explains how this relationship works.

FIGURE 1.35 The 64-bit memory throughput formula

FSB in MHz	(~~cycles~~/second)
X 8 bytes	(bytes/~~cycle~~)
throughput	(bytes/second)

Because the actual system clock speed is rarely mentioned in marketing literature, on packaging, or on store shelves for DDR and higher, you can use this advertised FSB frequency in your computations for DDR throughput. For example, with a 100 MHz clock and two operations per cycle, motherboard makers will market their boards as having an FSB of 200 MHz. Multiplying this effective rate by 8 bytes transferred per cycle, the data rate is 1,600 MBps. Because DDR made throughput a bit trickier to compute, the industry began using this final throughput figure to name the memory modules instead of the actual frequency, which was used when naming SDR modules. This makes the result seem many times better (and much more marketable), while it's really only twice (or so) as good, or close to it.

In this example, the module is referred to as PC1600, based on a throughput of 1,600 MBps. The chips that go into making PC1600 modules are named DDR200 for the effective FSB frequency of 200 MHz. Stated differently, the industry uses DDR200 memory chips to manufacture PC1600 memory modules.

Let's make sure that you grasp the relationship between the speed of the FSB and the name for the related chips as well as the relationship between the name of the chips (or the speed of the FSB) and the name of the modules. Consider an FSB of 400 MHz, meaning an actual clock signal of 200 MHz, by the way—the FSB is double the actual clock for DDR, remember. It should be clear that this motherboard requires modules populated with DDR400 chips and that you'll find such modules marketed and sold as PC3200.

Let's try another. What do you need for a motherboard that features a 333 MHz FSB (actual clock is 166 MHz)? Well, just using the 8:1 rule mentioned earlier, you might be on the lookout for a PC2667 module. Note, however, that sometimes the numbers have to be played with a bit to come up with the industry's marketing terms. You'll have an easier time finding PC2700 modules that are designed specifically for a motherboard like yours, with an FSB of 333 MHz. The label isn't always technically accurate, but round numbers sell better, perhaps. The important concept here is that if you find PC2700 modules and PC2667 modules, there's absolutely no difference; they both have a 2667 MBps throughput rate. Go for the best deal; just make sure that the memory manufacturer is reputable.

DDR2 SDRAM

Think of the *2* in *DDR2* as yet another multiplier of 2 in the SDRAM technology, using a lower peak voltage to keep power consumption down (1.8V vs. the 2.5V of DDR). Still double-pumping, DDR2, like DDR, uses both sweeps of the clock signal for data transfer. Internally, DDR2 further splits each clock pulse in two, doubling the number of operations it can perform per FSB clock cycle. Through enhancements in the electrical interface and buffers, as well as through adding off-chip drivers, DDR2 nominally produces four times the throughput that SDR is capable of producing.

Continuing the DDR example, DDR2, using a 100 MHz actual clock, transfers data in four operations per cycle (effective 400 MHz FSB) and still 8 bytes per operation, for a total of 3,200 MBps. Just as with DDR, chips for DDR2 are named based on the perceived frequency. In this case, you would be using DDR2-400 chips. DDR2 carries on the effective FSB frequency method for naming modules but cannot simply call them PC3200 modules because those already exist in the DDR world. DDR2 calls these modules PC2-3200. (Note the dash to keep the numeric components separate.)

As another example, it should make sense that PC2-5300 modules are populated with DDR2-667 chips. Recall that you might have to play with the numbers a bit. If you multiply the well-known FSB speed of 667 MHz by 8 to figure out what modules you need, you might go searching for PC2-5333 modules. You might find someone advertising such modules, but most compatible modules will be labeled PC2-5300 for the same marketability mentioned earlier. They both support 5,333 MBps of throughput.

DDR3 SDRAM

The next generation of memory devices was designed to roughly double the performance of DDR2 products. Based on the functionality and characteristics of DDR2's proposed successor, most informed consumers and some members of the industry surely assumed the forthcoming name would be DDR4. This was not to be, however, and DDR3 was born. This naming convention proved that the *2* in DDR2 was not meant to be a multiplier but instead a revision mark of sorts. Well, if DDR2 was the second version of DDR, then DDR3 is the third. *DDR3* is a memory type, designed to be twice as fast as the DDR2 memory, that operates with the same system clock speed. Just as DDR2 was required to lower power consumption to make up for higher frequencies, DDR3 must do the same. In fact, the peak voltage for DDR3 is only 1.5V.

The most commonly found range of actual clock speeds for DDR3 tends to be from 133 MHz at the low end to less than 300 MHz. Because double-pumping continues with DDR3, and because four operations occur at each wave crest (eight operations per cycle), this frequency range translates to common FSB implementations from 1,066 MHz to more than 2,000 MHz in DDR3 systems. These memory devices are named following the conventions established earlier. Therefore, if you buy a motherboard with a 1,600 MHz FSB, you know immediately that you need a memory module populated with DDR3-1600 chips, because the chips are always named for the FSB speed. Using the 8:1 module-to-chip/FSB naming rule, the modules that you need would be called PC3-12800, supporting a 12,800 MBps throughput.

The earliest DDR3 chips, however, were based on a 100 MHz actual clock signal, so we can build on our earlier example, which was also based on an actual clock rate of 100 MHz. With eight operations per cycle, the FSB on DDR3 motherboards is rated at 800 MHz, quite a lot of efficiency while still not needing to change the original clock with which our examples began. Applying the 8:1 rule again, the resulting RAM modules for this motherboard are called PC3-6400 and support a throughput of 6,400 MBps, carrying chips called DDR3-800, again named for the FSB speed.

DDR4 SDRAM

Continuing the inevitable march of technology, *DDR4* was the next iteration of SDRAM on the market. As you would expect, the speed of DDR4 is roughly double that of DDR3. DDR3 provided data rates of approximately 800 Mbps to 2,133 Mbps, whereas DDR4 ranges between 1,600 Mbps and 3,200 Mbps. DDR also runs at a lower voltage—1.2 volts. Finally, DDR4 can support more memory per module, up to 512 GB per chip. Realistically, though, no one produces motherboards (or RAM) that support that quantity. The largest you will see are 64 GB sticks.

DDR5 SDRAM

After a long wait, *DDR5* finally hit the market at the end of 2021. Intel's Alder Lake platform was the first to support it; AMD chips could support DDR5 in early 2022 with the Zen 4 release.

DDR5 doubles the speed of DDR4 to 6.4 Gbps, as is expected for a new memory standard. Improved power efficiency means it runs at 1.1 volts. DDR5 is also the first memory module to be available in up to 128 GB modules.

Exam Objective 3.3 states that test-takers should be able to compare and contrast DDR iterations. It's highly unlikely that you will need to know intimate details of each. What's more likely to be important are a few facts:

- Each version of DDR doubles the speed of the previous version.
- You might have noticed this in the discussion on DDR or in the real world, but some RAM is labeled "PC*X*" whereas others are "DDR*X*." The *X* will be the same within a version. In other words, PC4 RAM is DDR4 RAM.
- Motherboards will only physically support one type of DDR (more on that in the upcoming "Memory Packaging" section), such as DDR4 or DDR5. The motherboard documentation will tell you which type it supports.
- Motherboards will only work with certain speeds of RAM. It used to be that there was only one speed (or maybe two speeds) of RAM a motherboard would support, such as PC3-12800, based on its chipset. Now, motherboards are more flexible and can support different RAM speeds—but it's best practice to install only one speed in a system. And of course, faster is better. Look at the motherboard documentation to see what's supported. For example, a recent AORUS ELITE motherboard supports 23 different DDR5 speeds, from DDR5 5600 (PC5-44800) through DDR5 9200 (PC5-73600).

SRAM

Static random access memory (SRAM) doesn't require a refresh signal like DRAM does. The chips are more complex and are thus more expensive. However, they are considerably faster. DRAM access times can come in at 40 nanoseconds (ns) or more; SRAM has access times faster than 10ns. SRAM is classically used for cache memory.

ROM

ROM stands for read-only memory. It is called read-only because you could not write to the original form of this memory. Once information had been etched on a silicon chip and manufactured into the ROM package, the information couldn't be changed. Some form of ROM is normally used to store the computer's BIOS because this information normally does not change often.

The system ROM in the original IBM PC contained the power-on self-test (POST), BIOS, and cassette BASIC, a simple programming language for writing and running programs. Later, IBM computers and compatibles included everything but cassette BASIC. The system ROM enables the computer to "pull itself up by its bootstraps," or *boot* (find and start the operating system).

Through the years, different forms of ROM were developed that could be altered, later ones more easily than earlier ones. The first generation was the programmable ROM (PROM), which could be written to for the first time in the field using a special programming device, but then no more. You may liken this to the burning of a DVD-R.

The erasable PROM (EPROM) followed the PROM, and it could be erased using ultraviolet light and subsequently reprogrammed using the original programming device. These days, flash memory is a form of electronically erasable PROM (EEPROM). Of course, it does not require UV light to erase its contents, but rather a slightly higher than normal electrical pulse.

Although the names of these memory devices are different, they all contain ROM. Therefore, regardless which of these technologies is used to manufacture a BIOS chip, it's never incorrect to say that the result is a ROM chip.

Memory Packaging

The memory slots on a motherboard are designed for particular module form factors or styles. RAM historically evolved from form factors no longer seen for such applications, such as dual in-line package (DIP), single in-line memory module (SIMM), and single in-line pin package (SIPP). The most popular form factors for primary memory modules today are as follows:

- DIMM (dual in-line memory module)
- SODIMM (small outline dual in-line memory module)

Desktop computers will use DIMMs. Laptops and smaller devices require SODIMMs or smaller memory packaging. Other devices may use proprietary formats. So, in addition to coordinating the speed of the components, their form factor is an issue that must be addressed.

Exam Objective 3.3 states that you should be able to compare and contrast RAM form factors: DIMM and SODIMM.

DIMM

The most popular desktop PC memory package is known as a DIMM, which stands for dual in-line memory module. DIMMs are 64-bit memory modules that are used as a package for the SDRAM family: SDR, DDR, DDR2, DDR3, DDR4, and DDR5. The term *dual* refers to

the fact that, unlike their SIMM predecessors, DIMMs differentiate the functionality of the pins on one side of the module from the corresponding pins on the other side. With 84 pins per side, this makes 168 independent pins on each standard SDR module.

The DIMMs used for DDR memory are all the same length (5.25″ or 133.35mm), but they differ in the number of pins they have as well as where the key is in the module. Some will have built-in heat sinks, and newer modules may also have RGB lighting. Because of the position of the key, you can't put one type of DIMM into a different type of slot. Table 1.4 summarizes some key differences between the types of DDR we've introduced in this chapter.

Figure 1.36 shows SDR (on the left) and the DDR family in order from DDR to DDR5, left to right.

TABLE 1.4 DDR DIMM characteristics

Characteristic	DDR	DDR2	DDR3	DDR4	DDR5
Pins	184	240	240	288	288
Max memory	1 GB	8 GB	32 GB	64 GB	512 GB
Channels	1	1	1	1	2
Voltage	2.5 v	1.8 v	1.5 v	1.2 v	1.1 v

FIGURE 1.36 SDR and DDR DIMMs

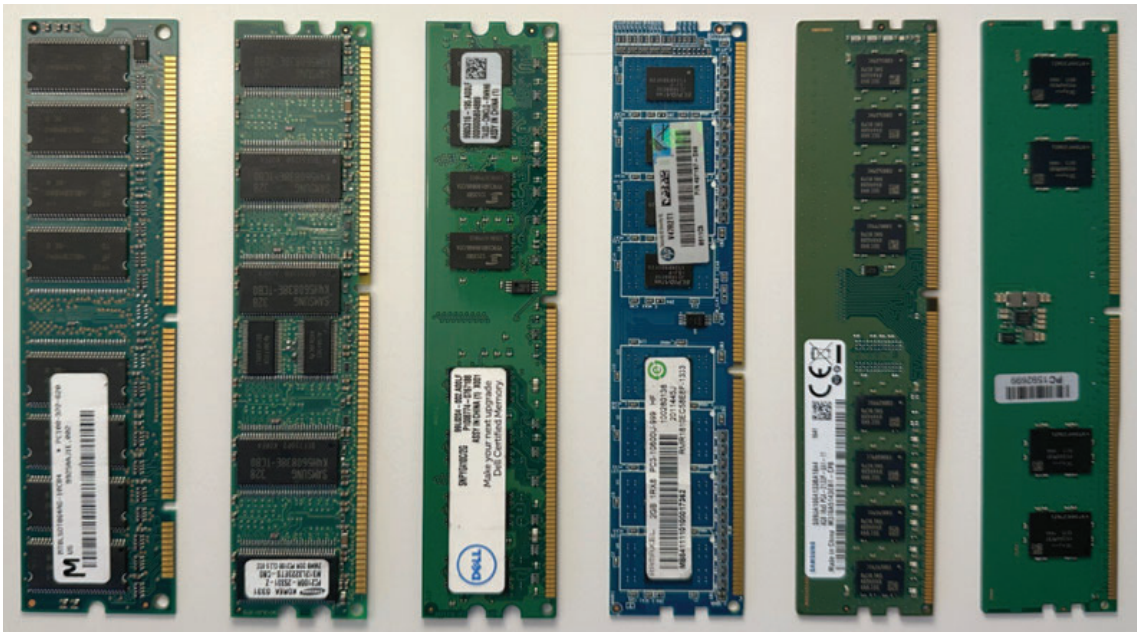

Inserting and Removing Memory Modules

The original single in-line memory modules had to be inserted into their slots at a 45° angle. The installer then had to apply slight pressure as the module was maneuvered upright at a 90° angle to the motherboard, where a locking mechanism would grip the module and prevent it from returning to its 45° position. This procedure created a pressure that reinforced the contact of the module with its slot. Releasing the clips on either end of the module unlocked it and allowed it to return to 45°, where it could be removed.

DIMM slots, by comparison, have no spring action. DIMMs are inserted straight into the slot with the locking tabs pulled away from the module. The locking tabs are at either end of the module, and they automatically snap into place, securing the module. Pulling the tabs away from the module releases the module from the slot, allowing it to be effortlessly removed.

SODIMM

Laptop computers and other computers that require much smaller components don't use standard RAM packages, such as DIMMs. Instead, they call for a much smaller memory form factor, such as a *small outline DIMM (SODIMM)*. SODIMMs are a little more than half as wide as DIMMs, at 2.75″ (69.6mm), and have fewer pins (outlined in Table 1.5). Each version is keyed differently as well, to prevent putting the wrong chip into a motherboard. SODIMMs use the same DDR electronic specifications as DIMMs, so the number of channels, voltage, and maximum amount of memory are similar.

TABLE 1.5 SODIMM pin counts

Version	SDR	DDR	DDR2	DDR3	DDR4	DDR5
Pins	144	200	200	204	260	262

Figure 1.37 shows three SODIMMs, DDR3 through DDR5, from top to bottom. Figure 1.38 shows a DDR5 SODIMM and a DDR5 DIMM for comparison.

Understanding Cooling Systems

It's a basic concept of physics: electronic components turn electricity into work and heat. The excess heat must be dissipated or it will shorten the life of the components. In some cases (like with the CPU), the component will produce so much heat that it can destroy itself in a matter of seconds if there is not some way to remove this extra heat.

FIGURE 1.37 SODIMMs

FIGURE 1.38 DDR5 SODIMM and DIMM

Air-cooling methods are used to cool the internal components of most PCs. With air cooling, the movement of air removes the heat from the component. Sometimes, large blocks of metal called *heat sinks* are attached to a heat-producing component in order to dissipate the heat more rapidly.

The following cooling components are included in Exam Objective 3.5: fans, heat sink, thermal paste/pads, and liquid.

Case Cooling and Fans

When you turn on a computer, you will often hear lots of whirring. Most of this noise is coming from the various fans inside the computer, although older hard disk drives (HDDs) make noise too. Fans provide airflow within the computer.

Most PCs have at a minimum three types of fans:

Front Intake Fan This fan is used to bring fresh, cool air into the computer for cooling purposes. Many newer cases have two or even three front intake fans and may also have one on the top of the case as well.

Rear Exhaust Fan This fan is used to take hot air out of the case. Newer systems may have multiple exhaust fans as well.

Power Supply Exhaust Fan This fan is usually found at the back of the power supply, and it is used to cool the power supply. In addition, this fan draws air from inside the case into vents in the power supply. This pulls hot air through the power supply so that it can be blown out of the case. The front intake fan assists with this airflow. The rear exhaust fan supplements the power supply fan to achieve the same result outside of the power supply.

Ideally, the airflow inside a computer should resemble what is shown in Figure 1.39, where the back of the chassis is shown on the left in the image.

Note that you must pay attention to the orientation of the power supply's airflow. If the power supply fan is an exhaust fan, as assumed in this discussion, the front and rear fans will match their earlier descriptions: front, intake; rear, exhaust. If you run across a power supply that has an intake fan, the orientation of the supplemental chassis fans should be reversed as well. The rear chassis fan(s) should always be installed in the same orientation the power supply fan runs to avoid creating a small airflow circuit that circumvents the cross-flow of air throughout the case. The front chassis fan and the rear fans should always be installed in reverse orientation to avoid having them fight against each other and thereby reduce the internal airflow. Reversing supplemental chassis fans is usually no harder than removing four screws and flipping the fan. Sometimes, the fan might just snap out, flip, and then snap back in, depending on the way it is rigged up.

FIGURE 1.39 System unit airflow

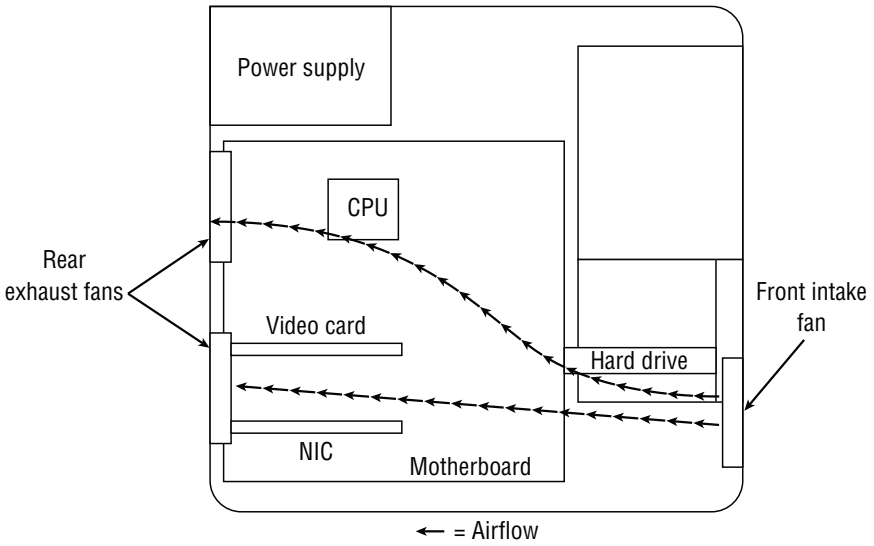

Motherboard Fan Power Connectors

It's important to be aware of the two main types of fan connections found on today's motherboards. One of these connectors has only three connections, while the other has four. The fan connectors and motherboard headers are interchangeable between the two pinouts, but if a chassis fan has four conductors, it's a sign that it's calling for connectivity to an extra +5VDC (volts direct current) connection that the most common 3-pin header doesn't offer. A more rare 3-pin chassis-fan connector features a +12VDC power connection for heavier-duty fans and a rotation pin used as an input to the motherboard for sensing the speed of the fan.

Four-pin CPU connections place the ground and power connections in pins 1 and 2, respectively, so that 2-pin connectors can be used to power older fans. The 4-pin header also offers a tachometer input signal from the fan on pin 3 so that the speed of the fan can be monitored by the BIOS and other utilities. Look for markings such as *CPU FAN IN* to identify this function. Pin 4 might be labeled *CPU FAN PWM* to denote the pulse-width modulation that can be used to send a signal to the fan to control its speed. This is the function lost when a 3-pin connector is placed in the correct position on a 4-pin header. Four-pin chassis-fan connectors can share the tachometer function but replace the speed control function with the extra 5V mentioned earlier.

Other power connections and types will be covered in Volume 1 Chapter 2, including the Molex connector, which can be used to power chassis and CPU fans using an adapter or the built-in connector on mostly older fans manufactured before the motherboard connectors were standardized. Figure 1.40 shows two 3-pin chassis-fan headers on a motherboard.

FIGURE 1.40 3-pin chassis-fan headers

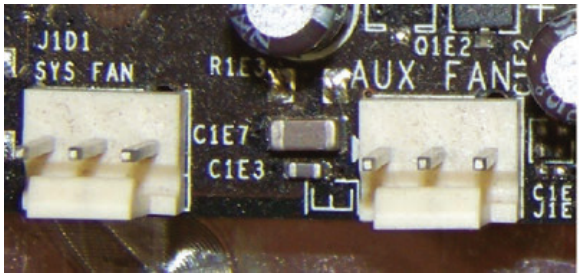

Figure 1.41 shows a 4-pin CPU fan header with an approaching 3-pin connector from the fan. Note that the keying tab is lined up with the same three pins it's lined up with in the 3-pin connectors.

FIGURE 1.41 A 4-pin CPU fan header

This physical aspect and the similar pin functions are what make these connectors interchangeable, provided that the header's function matches the role of the fan being connected. Figure 1.42 shows the resulting unused pin on the 4-pin header. Again, controlling the fan's speed is not supported in this configuration.

FIGURE 1.42 Position of a 3-pin connector on a 4-pin header

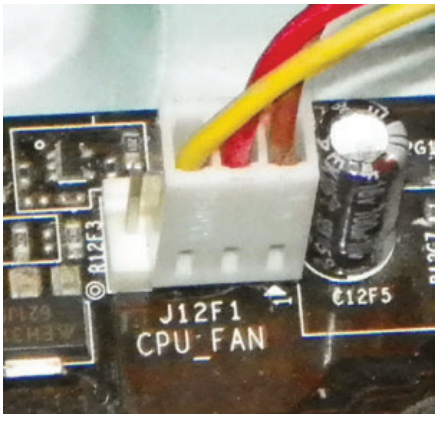

CPU Cooling

Probably the greatest challenge in cooling is the computer's CPU. It is the component that generates the most heat in a computer (aside from some pretty insane GPUs out there). As a matter of fact, if a modern processor isn't actively cooled all of the time, it will generate enough heat to burn itself up in seconds. That's why most motherboards have an internal CPU heat sensor and a CPU fan sensor, the connector for which is labeled CPU_FAN on the motherboard (as in Figure 1.41 and Figure 1.42). If no cooling fan is active, these devices will shut down the computer before damage occurs.

There are multiple CPU cooling methods, but the two most common are air cooling and liquid cooling.

Air Cooling

The parts inside most computers are cooled by air moving through the case. The CPU is no exception. However, because of the large amount of heat produced, the CPU must have (proportionately) the largest surface area exposed to the moving air in the case. Therefore, the heat sinks on the CPU are the largest of any inside the computer.

The CPU fan often blows air down through the body of the heat sink to force the heat into the ambient internal air where it can join the airflow circuit for removal from the case. However, in some cases, you might find that the heat sink extends up farther, using radiator-type fins, and the fan is placed at a right angle and to the side of the heat sink. This design moves the heat away from the heat sink immediately instead of pushing the air down through the heat sink. CPU fans can be purchased that have an adjustable rheostat to allow you to dial in as little airflow as you need, aiding in noise reduction but potentially leading to accidental overheating.

It should be noted that the highest-performing air CPU coolers use copper plates in direct contact with the CPU. They also use high-speed and high-CFM cooling fans to dissipate the heat produced by the processor. CFM is short for cubic feet per minute, an airflow measurement of the volume of air that passes by a stationary object per minute. Figure 1.43 shows a newer, large heat sink with a fan in the center. In the picture it can be tough to gauge size—this unit is about 6″ across! (And to be fair, this heat sink should have a second fan on one of the sides, but with the second fan the heatsink wouldn't fit into the author's case—the RAM was in the way.)

Many new CPU heat sinks use tubing to transfer heat away from the CPU. With any cooling system, the more surface area exposed to the cooling method, the better the cooling. Plus, the heat pipes can be used to transfer heat to a location away from the heat source before cooling. This is especially useful in cases where the form factor is small and with laptops, where open space is limited.

With advanced heat sinks and CPU cooling methods like this, it is important to improve the thermal transfer efficiency as much as possible. To that end, cooling engineers came up with a glue-like compound that helps to bridge the extremely small gaps between the CPU and the heat sink, which avoids superheated pockets of air that can lead to focal damage of the CPU. This product is known as thermal transfer compound, or simply thermal compound (alternatively, thermal grease or *thermal paste*), and it can be bought in small tubes. Single-use tubes are also available and alleviate the guesswork involved with how much you should apply. Watch out, though; this stuff makes quite a mess and doesn't want

to come off your fingers very easily. An alternative to the paste is a small *thermal pad*, which provides heat conductivity between the processor and the heat sink. Figure 1.44 shows a tube of thermal paste and a thermal pad (the pad needs to be cut to size).

FIGURE 1.43 Large heat sink and fan

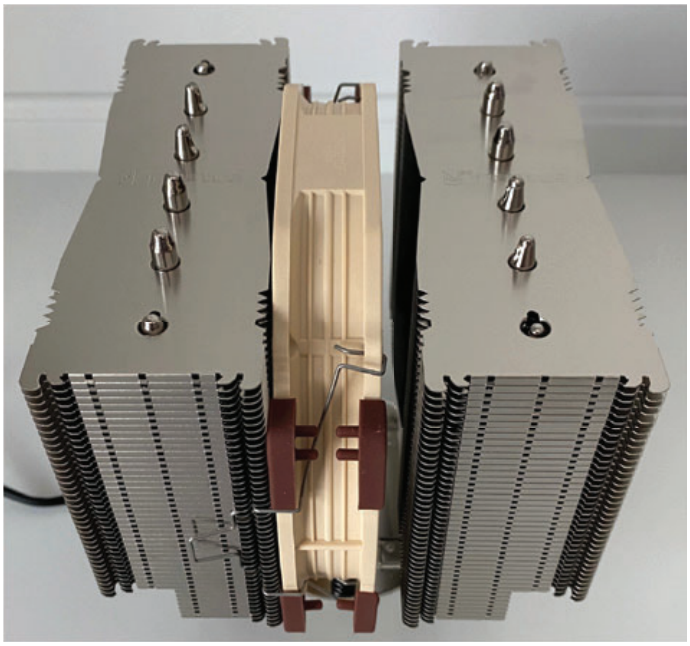

FIGURE 1.44 Thermal paste and a thermal pad

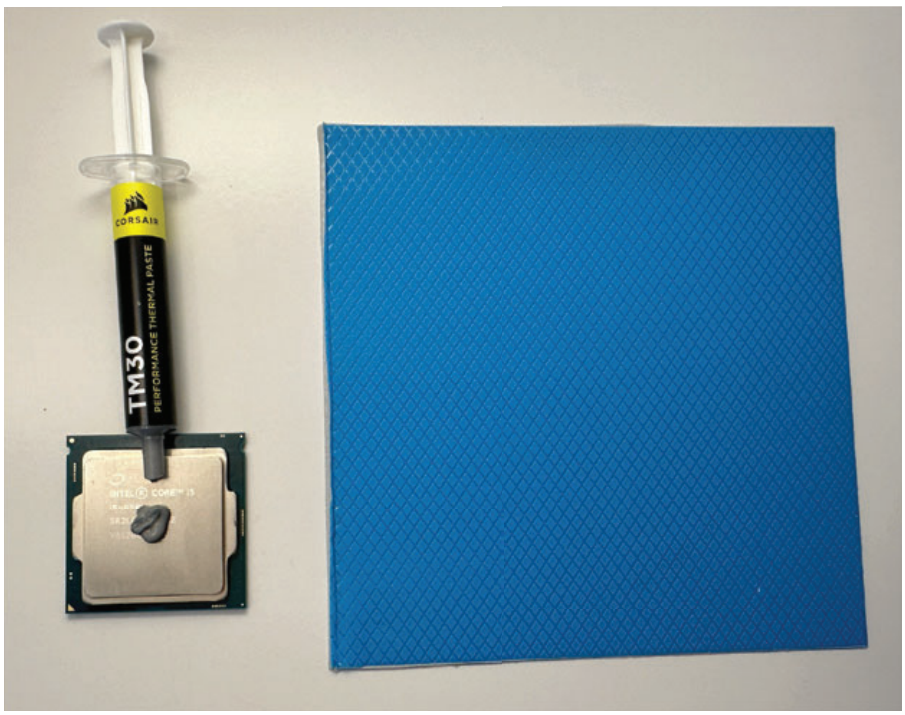

Apply the compound by placing a bead in the center of the heat sink, not on the CPU, because some heat sinks don't cover the entire CPU package. That might sound like a problem, but some CPUs don't have heat-producing components all the way out to the edges. Some CPUs even have a raised area directly over the silicon die within the packaging, resulting in a smaller contact area between the components. You should apply less than you think you need because the pressure of attaching the heat sink to the CPU will spread the compound across the entire surface in a very thin layer. It's advisable to use a clean, lint-free applicator of your choosing to spread the compound around a bit as well, just to get the spreading started. You don't need to concern yourself with spreading it too thoroughly or too neatly because the pressure applied during attachment will equalize the compound quite well. During attachment, watch for oozing compound around the edges, clean it off immediately, and use less next time.

Improving and Maintaining CPU Cooling

In addition to using thermal compound, you can enhance the cooling efficiency of a CPU heat sink by lapping the heat sink, which smoothens the mating surface using a very fine sanding element, about 1,000 grit in the finishing stage. Some vendors of the more expensive heat sinks will offer this service as an add-on.

If your CPU has been in service for an extended period of time, perhaps three years or more, it is a smart idea to remove the heat sink and old thermal compound and then apply fresh thermal compound and reattach the heat sink. Be careful, though; if your thermal paste has already turned into thermal "glue," you can wrench the processor right out of the socket, even with the release mechanism locked in place. Invariably, this damages the pins on the chip. Try running the computer for a couple of minutes to warm the paste and then try removing the heat sink again.

Counterintuitively, perhaps, you can remove a sealed heat sink from the processor by gently rotating the heat sink to break the paste's seal. Again, this can be made easier with heat. If the CPU has risen in the socket already, however, rotating the heat sink would be an extremely bad idea. Sometimes, after you realize that the CPU has risen a bit and that you need to release the mechanism holding it in to reseat it, you find that the release arm is not accessible with the heat sink in place. This is an unfortunate predicament that will present plenty of opportunity to learn.

If you've ever installed a brand-new heat sink onto a CPU, you've most likely used thermal compound or the thermal compound patch that was already applied to the heat sink for you. If your new heat sink has a patch of thermal compound pre-applied, don't add more. If you ever remove the heat sink, don't try to reuse the patch or any other form of thermal compound. Clean it all off and start fresh.

Liquid Cooling

Liquid cooling is a technology whereby a special water block is used to conduct heat away from the processor (as well as from the chipset). Water is circulated through this block to a radiator, where it is cooled.

The theory is that you could achieve better cooling performance through the use of liquid cooling. For the most part, this is true. However, with traditional cooling methods (which use air and water), the lowest temperature you can achieve is room temperature. Plus, with liquid cooling, the pump is submerged in the coolant (generally speaking), so as it works, it produces heat, which adds to the overall liquid temperature.

The main benefit to liquid cooling is silence. Only one fan is needed: the fan on the radiator to cool the water. So, a liquid-cooled system can run extremely quietly.

There are two major classifications of liquid cooling systems in use with PCs today: all-in-one (AIO) coolers and custom loop systems. AIO systems are relatively easy to install—they require about as much effort as a heat sink and fan—and are comparably priced to similarly effective air systems. Figure 1.45 shows an example from Corsair, with the pump in front and the fans behind it, attached to the radiator.

FIGURE 1.45 AIO liquid cooling system

AIO systems come in three common sizes: 120mm (with one fan, and the most common), 240mm (two fans, for overclocked components), and 360mm (three fans, for high-end multicore overclocked components). Options with RGB lighting are readily available if that's the style you want. Custom loop systems can quickly become complex and expensive, but many hardcore gamers swear by their performance. The components are essentially the same as those in an AIO system—there's a radiator, pump, fans, some tubes, and liquid. However, each part is purchased separately, and some assembly is required.

Water cooling systems often require more room inside the case than heat sinks, as well as special headers on the motherboard (or in the case) to support the fans. As always, check your documentation to be sure that your other hardware is compatible!

Cooling Other Components

As noted previously, the CPU is the most important component to cool. Most other components receive adequate cooling from the case fans, provided they are working properly and not covered in a thick layer of dust. There are, however, cooling systems for other internal components as well. Most of them involve the use of heat sinks and/or fans.

Video Card Cooling

Modern video cards run almost as hot as CPUs and will come with cooling mechanisms built-in to their packaging. Most of the time this is a fan or series of fans covering heat sinks. We will show an image of a video card with cooling fans in Volume 1 Chapter 2.

Memory Cooling

If you are going to start overclocking your computer, you will want to do everything in your power to cool all of its components, and that includes the memory.

There are two methods of cooling memory: passive and active. The passive memory cooling method just uses the ambient case airflow to cool the memory through the use of enhanced heat dissipation. For this, you can buy either heat sinks or special "for memory chips only" devices known as *heat spreaders*. These are special aluminum or copper housings that wrap around memory chips and conduct the heat away from them.

Active cooling, on the other hand, usually involves forcing some kind of cooling medium (air or water) around the RAM chips themselves or around their heat sinks. Most often, active cooling methods are just high-speed fans directing air right over a set of heat spreaders.

Hard Drive Cooling

There are both active and passive cooling devices for hard drives. Most common, however, is the active cooling bay. You install a hard drive in a special device that fits into a standard hard drive expansion bay. This device contains fans that draw in cool air over the hard drive, thus cooling it. Figure 1.46 shows an example of one of these active hard drive coolers, with two fans, mounted to a 3.5″ hard drive. As you might suspect, you can also get heat sinks for hard drives.

FIGURE 1.46 An active hard disk cooler

Summary

In this chapter, we took a tour of the key internal system components of a PC. You learned primarily about the "big three," which are the motherboard, processor, and memory. Included in the motherboard discussion were form factors, connector types, such as PCIe, M.2, and headers, motherboard and CPU compatibility, BIOS/UEFI settings, encryption, and the CMOS battery. CPU topics included features such as compatibility, architecture, multithreading, and virtualization. With RAM you learned about different types (SODIMMs and DIMMs, DDR2, DDR3, DDR4, and DDR5) and concepts such as single-, dual-, triple-, and quad-channel, error correcting, and parity.

Finally, the chapter ended with cooling systems, which keep the components from damaging themselves with excess heat. This chapter laid the foundation for the rest of the book, including the next few chapters on additional hardware components.

Exam Essentials

Know three form factors of system boards. Know the characteristics of and differences between ATX, microATX, and ITX motherboards.

Know the components of a motherboard. Be able to describe, identify, and replace (where applicable) motherboard components, such as expansion slots, memory slots, processor sockets, BIOS/UEFI (firmware), and CMOS batteries.

Be able to identify and differentiate motherboard connector types. Understand the differences between PCI, PCIe, SATA, eSATA, and M.2 connectors, as well as power connectors and headers.

Understand core concepts of motherboard compatibility. Know that Intel and AMD chips use different sockets and therefore are incompatible with each other. Also know differences between multisocket, desktop, and laptop motherboards.

Understand CPU architecture. Know the differences between x64, x86, and ARM processors, implications of single-core versus multicore CPUs, and multithreading and virtualization support.

Know what the BIOS/UEFI is responsible for. The BIOS/UEFI controls boot options, fan speeds, temperature monitoring, USB permissions, and security options such as the BIOS password, boot password, TPM, HSM, and Secure Boot.

Understand the purposes and characteristics of memory. Know about the characteristics that set the various types of memory apart from one another. This includes the actual types of memory, such as DRAM (which includes several varieties), SRAM, ROM, and CMOS, as well as memory packaging, such as DIMMs and SODIMMs. Also have a firm understanding of the different levels of cache memory as well as its purpose in general.

Know how to replace RAM, given a scenario. RAM must match the motherboard in form factor and in speed. For example, some motherboards will only accept DDR4 or DDR5 memory. The speed should be compatible with the motherboard, as indicated in the documentation.

Understand the purposes and characteristics of cooling systems. Know the different ways internal components can be cooled and how overheating can be prevented. Most cooling systems involve a fan, a heat sink, or a combination of the two, or liquid cooling and a fan. Heat sinks and fans are often attached with thermal paste or a thermal pad.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. Which computer component contains all the circuitry necessary for other components or devices to communicate with one another?
   - **A.** Motherboard
   - **B.** CPU
   - **C.** RAM
   - **D.** Expansion bus

2. You are told by a technician over the phone that you need to bring replacement DDR4 SDRAM memory. What type of packaging is used for DDR4 SDRAM memory?
   - **A.** 224-pin DIMM
   - **B.** 240-pin DIMM
   - **C.** 288-pin DIMM
   - **D.** 296-pin DIMM

3. You need to replace memory in a server that is using ECC RAM. Which of the following statements regarding the replacement RAM is true?
   - **A.** Installing non-ECC RAM will work fine, but the new RAM won't have error-correcting.
   - **B.** You can only install ECC RAM for it to work properly.
   - **C.** You can install non-ECC RAM and it will adapt for error-correction.
   - **D.** Installing non-ECC RAM will work fine as long as the new RAM includes a parity bit.

4. A client wants you to build a new PC for her, with a smaller case and lower power requirements. When selecting a motherboard, which form factor should you choose for a smaller size and lower power consumption?
   - **A.** ATX
   - **B.** AT
   - **C.** DTX
   - **D.** ITX

5. A motherboard failed on your desktop PC, which has an Intel Core i5-13600K processor. When you're looking for a replacement motherboard, which CPU socket does it need to have?
   - **A.** LGA 1366
   - **B.** LGA 1150
   - **C.** LGA 1200
   - **D.** LGA 1700

6. You have just replaced a processor in a computer and now need to add a cooling mechanism. What should you use to attach the cooling system to the processor?
   - **A.** Heat sink
   - **B.** Thermal paste
   - **C.** Fan
   - **D.** Superglue

7. A technician asks you how to get a computer's processor to run faster than it currently does. What is he trying to achieve?
   - **A.** Hyperthreading
   - **B.** Overclocking
   - **C.** Virtualization
   - **D.** Multicore support

8. You are assembling a new computer and are looking at new cooling systems. Which of the following cooling systems requires the use of a pump?
   - **A.** Liquid
   - **B.** Thermal paste
   - **C.** Heat sink
   - **D.** Heat sink plus fan

9. Which of the following types of processors will likely generally be preferred for mobile devices?
   - **A.** x64
   - **B.** x86
   - **C.** ARM
   - **D.** LGA

10. You press the front power button on a computer and the system boots. Later, you press it briefly and the system hibernates. When you press it again, the system resumes. You press and hold the button and the system shuts down. What is this feature called?
    - **A.** Programmable power
    - **B.** Soft power
    - **C.** Relay power
    - **D.** Hot power

11. You are training new technicians to install RAM and pointing out differences between packages. Which of the following are the numbers of pins that can be found on DIMM modules used in desktop motherboards? (Choose two.)
    - **A.** 180
    - **B.** 184

C. 200

D. 204

E. 232

F. 240

12. You are installing a new computer, and the user wants to use virtualization. Which two hardware components need to support virtual technology for this to work properly? (Choose two.)

A. RAM

B. Motherboard

C. CPU

D. BIOS

13. You find out that a disgruntled ex-employee's computer has a boot password that must be entered before the operating system is ever loaded. There is also a password preventing your access to the BIOS utility. Which of the following motherboard components can most likely be used to return the computer to a state that will allow you to boot the system without knowing the password?

A. Cable header

B. Power reset connector

C. Toggle switch

D. Jumper

14. Your Core i5 fan has a 4-pin connector, but your motherboard only has a single 3-pin header with the CPU_FAN label. Which of the following will be the easiest solution to get the necessary cooling for your CPU?

A. Plug the 4-pin connector into the 3-pin header.

B. Buy a 4-pin to 3-pin adapter.

C. Leave the plug disconnected and use only the heat sink.

D. Add an extra chassis fan.

15. You are installing a new video card into a PCIe slot. What is the combined total throughput of a PCIe 4.0 x16 slot?

A. 500 MBps

B. 1 GBps

C. 64 GBps

D. 128 GBps

16. Which of the following allows you to perform the most complete restart of the computer without removing power?

A. Start ➢ Restart

B. Start ➢ Hibernate

C. Reset button

D. Power button

17. You have just found out that a new UEFI upgrade exists for your computer, and you obtain a copy on a flash drive to install it. Which of the following is most helpful when flashing the UEFI on a desktop computer system?

A. DVD-ROM

B. Uninterruptable power supply

C. Internet connection

D. Windows administrator password

18. The 32 GB of DDR4 memory has failed in a client's desktop computer. The client requests that you upgrade the system for better performance when replacing the RAM. Which of the following is the best option?

A. Install two new 32 GB DDR4 RAM modules.

B. Install two new 128 GB DDR4 RAM modules.

C. Install two new 32 GB DDR5 RAM modules.

D. Install two new 128 GB DDR5 RAM modules.

19. You are replacing an HP laptop's motherboard. Which type of motherboard should you use?

A. ATX

B. AT

C. One for the specific model

D. mini-ITX

20. You need to install a new video card into a desktop computer. What type of expansion slot is preferred today for high-performance graphics adapters?

A. AGP

B. PCIe

C. PCI

D. SATA

Performance-Based Question 1

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answer compares to mine, refer to Appendix B.

You have been asked to remove a dual in-line memory module and insert one with a larger capacity in its place. Describe the process for doing so.

Performance-Based Question 2

Identify the component each arrow points to in the following image of an ATX motherboard.

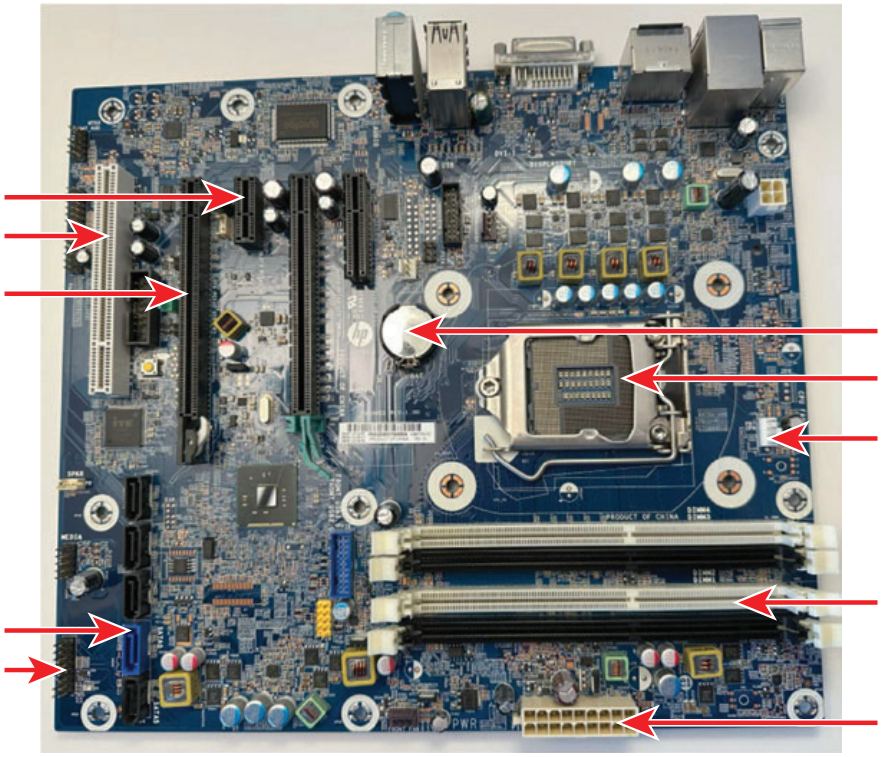

Chapter 2

Expansion Cards, Storage Devices, and Power Supplies

THE FOLLOWING COMPTIA A+ 220-1201 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

✓ **3.0 Hardware**

✓ **3.4 Compare and contrast storage devices.**

- Hard drives
  - Spindle speeds
  - Form factors
    - 2.5-inch
    - 3.5-inch
- Solid-state drives
  - Communications interfaces
    - Non-volatile Memory Express (NVMe)
    - SATA
    - Peripheral Component Interconnect Express (PCIe)
    - Serial Attached SCSI [Small Computer System Interface] (SAS)
  - Form factors
    - M.2
    - Mini-serial Advanced Technology Attachment (mSATA)
- Drive configurations
  - Redundant Array of Independent Disks (RAID) 0, 1, 5, 6, 10

- Removable storage
  - Flash drives
  - Memory cards
- Optical drives

✓ **3.5 Given a scenario, install and configure motherboards, central processing units (CPUs), and add-on cards.**

- Expansion cards
  - Sound card
  - Video card
  - Capture card
  - Network interface card

✓ **3.6 Given a scenario, install the appropriate power supply.**

- Input 110-120 VAC vs. 220-240 VAC
- Output 3.3V vs. 5V vs. 12V
- 20+4 pin motherboard connector
- Redundant power supply
- Modular power supply
- Wattage rating
- Energy efficiency

As a PC technician, you need to know quite a bit about hardware. Given the importance and magnitude of this knowledge, the best way to approach learning about it is in sections. The first chapter introduced the topic via the primary core components, and this chapter follows up where it left off. First, we will look at adding functionality by plugging expansion cards into the motherboard. Then, we will focus on storage devices that hold data persistently—that is, they don't require power to maintain data like RAM does. Finally, we will end the chapter by looking at the simple-looking but potentially dangerous box that gives the components the energy they need—the power supply.

Installing and Configuring Expansion Cards

An *expansion card* (also known as an *adapter card*) is a circuit board that you install into a computer to increase the capabilities of that computer. Expansion cards come in varying formats for different uses, but the important thing to note is that no matter what function a card has, the card being installed must match the bus type of the motherboard into which it is being installed. For example, you can install a PCIe network card into a PCIe expansion slot only.

For today's integrated components (those built into the motherboard), you might not need an adapter to achieve the related services. Each adapter card, integrated or not, needs a driver—a software program that lets the operating system talk to the hardware—to make the device function with the operating system. Most motherboard manufacturers supply drivers with their motherboards, typically on a flash drive, that contain all the device drivers needed to get the built-in electronics recognized by the operating system. Drivers are also almost always available on the device manufacturer's website. Execution of the driver's setup program generally results in all components working properly. We'll show you how to install and update drivers in the upcoming "Adapter Configuration" section.

The following are the four most common categories of expansion cards installed today:

- Video
- Multimedia
- Network interface
- Input/output

Let's take a quick look at each of these card types, their functions, and what some of them look like.

A+ Exam Objective 3.5, "Given a scenario, install and configure motherboards, central processing units (CPUs), and add-on cards," wants you to be familiar with sound cards, video cards, capture cards, and network interface cards.

Video

A *video card* (sometimes called a *graphics card*) is the expansion card that you put into a computer to allow the computer to present information on some kind of display, typically a monitor or a projector. A video card is also responsible for converting the data sent to it by the CPU into the pixels, addresses, and other items required for display. Video cards can include dedicated chips to perform some of these functions, thus accelerating the speed of display.

You will encounter two classes of video cards: onboard cards and add-on cards. Onboard (or integrated) cards are built into the motherboard. As mentioned earlier, you need to install a device driver to get them to work properly, but those often come packaged with the motherboard itself. The upside to an integrated card is that it frees up an expansion slot. The manufacturer can either leave the slot open or design the motherboard and/or case to be smaller. One downside is that if the video card fails, you need a new motherboard, or you can install an add-on card. A second downside is that the onboard video cards aren't typically high-end. Onboard cards generally share system memory with the processor, which limits the quality of graphics one can produce. If the user wants great graphics from a powerful video card, then an add-on card is almost always the way to go. For example, serious gamers will always insist on a separate video card.

Add-on cards will have their own dedicated *graphics processing unit (GPU)*, which handles all of the video processing and frees up CPU capacity. In addition, they will have their own dedicated memory, called *video random-access memory (VRAM)*, which frees up system RAM. PCIe is the preferred expansion slot type. You might be able to find the rare, outdated motherboard that still offers a legacy AGP slot, and you might see some cheap PCI video cards, but they are uncommon. The technology on which PCIe was designed performs better for video than those on which AGP and PCI are based. Figure 2.1 shows an example of a PCIe x16 video card. The video card pictured is 10.6″ (270mm) long and takes up quite a bit of space inside the case. Most cards today have built-in fans like this one does to reduce the chance of overheating.

There is an extensive range of video cards available today on the market. For everyday usage, cards with 2–4 GB of video memory are inexpensive and will do the trick. For gamers, high-end cards with a minimum of 12–16 GB GDDR5 are recommended. Of course, over the lifespan of this book, that number is sure to increase. (As of this writing, cards with 24 GB GDDR6 are available.) High-end video cards can easily cost several thousand dollars.

FIGURE 2.1 A PCIe video expansion card

The main two standards for video cards are the NVIDIA GeForce series and the AMD Radeon (formerly ATI Radeon) line. Gamers will debate the pros and cons of each platform but know that you can get a range of performance, from good to phenomenal, from either one. When looking for a card, know how much memory is wanted or needed and how many and which types of video ports (such as HDMI or DisplayPort) are available. We will talk more about the pros and cons of several video connectors in Volume 1 Chapter 3, "Peripherals, Cables, and Connectors."

Space May Be a Challenge

Older video cards were shorter than current high-end models, often only 4–6″ long. When manufacturers started creating faster GPUs and adding memory, cards got longer. Once cards got to about 12″ long, some of them would no longer fit into standard ATX cases without running into other components.

Faster GPUs and more memory take up space and generate heat, though, which requires larger cooling systems. So instead of going longer, video card manufacturers have gone wider. The EVGA GeForce RTX 2070 Super shown in Figure 2.1 is a dual-slot card, meaning it takes up the equivalent of two back-panel slots in a standard case. Current high-end models are quad slot. Because of their width, they might preclude the use of other expansion slots on a motherboard.

When looking at buying new video card, check the size specifications and make sure it will fit with the case and motherboard you're planning to use.

Multimedia

The most basic and prolific multimedia adapter is the sound card. Video capture cards also offer multimedia experiences but are less common than sound cards.

Sound Cards

Just as there are devices to convert computer signals into printouts and video information, there are devices to convert those signals into sound. These devices are known as *sound cards*. Although sound cards started out as pluggable adapters, this functionality is one of the most common integrated technologies found on motherboards today. A sound card typically has small, round 1/8″ (3.55mm) jacks on the back of it for connecting microphones, headphones, and speakers as well as other sound equipment. Older sound cards used a DA15 game port, which could be used for either joysticks or Musical Instrument Digital Interface (MIDI) controllers. Figure 2.2 shows an example of a sound card with five 3.55mm jacks and a DA15 game port. The audio jacks will be color-coded to identify their function. We know that this book isn't printed in color, so trust us when we say that from left to right, the ports are:

- **Yellow:** Digital output
- **Light blue:** Line level stereo input, such as from a sound mixer or other amplified audio equipment
- **Pink:** Microphone input
- **Lime green:** Stereo output or headphones
- **Black:** Surround sound left and right output

Some sound cards might also have an orange jack for center channel and subwoofer output, and a gray hack for surround sound rear left and right output. Other cards might have fewer ports, such as just light blue (line in), pink (microphone), and lime green (audio output).

FIGURE 2.2 A classic sound card

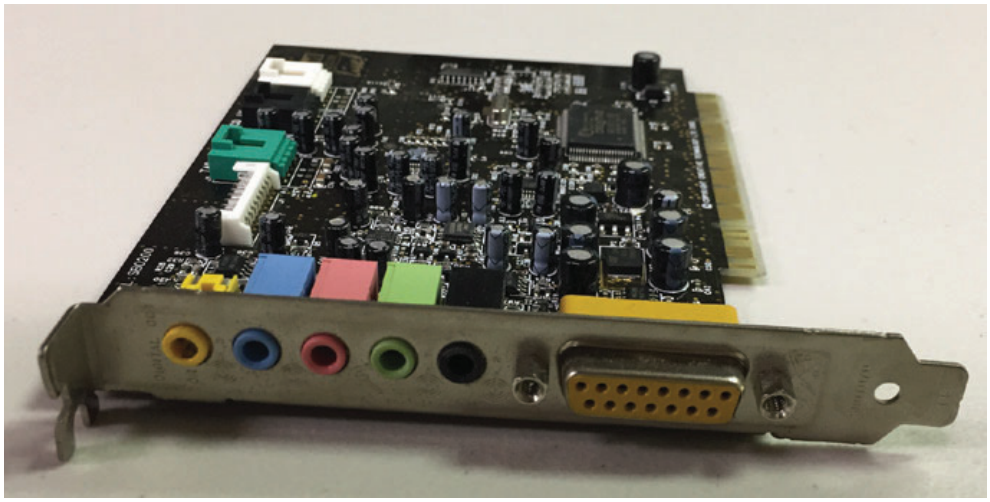

In our section on video cards, we noted that integrated cards have inferior performance to add-on ones, and though the same holds true for sound cards, the difference isn't quite as drastic. Many of today's motherboards come equipped with 5.1 or 7.1 analog or digital audio and support other surround sound formats as well. For everyday users and even many gamers, integrated audio is fine.

For users who need extra juice, such as those who produce movies or videos or do other audio/video (A/V) editing, a specialized add-on sound card is a must. Very good quality sound cards can be found for under $100, compared with cheaper models around $20—there isn't a huge difference in price as there is with CPUs and video cards. Look for a card with a higher sampling rate (measured in kilohertz [kHz]) and higher signal-to-noise ratio (measured in decibels [dB]). The de facto standard for sound cards is the Sound Blaster brand. Although other brands exist, they will often tout "Sound Blaster compatibility" in their advertising to show that they are legit.

Integrated motherboard audio will typically have about a 90 kHz sampling rate and 85–90 dB signal-to-noise ratio (SNR or S/N). A good add-on sound card can provide sampling of 190 kHz or more and an SNR of 115 dB or higher. The difference is noticeable, especially in professional productions that may be played at loud volumes.

In addition to audio output, many A/V editors will require the ability to input custom music from an electronic musical keyboard or other device. A term you will hear in relation to this is the MIDI standard. MIDI connections are almost always made via USB today, although you may occasionally see a round 5-pin MIDI connector. Look for the term *MIDI compatible* if attaching a digital musical device is needed.

Video Capture Cards

A *video capture card* is a stand-alone add-on card often used to save a video stream to the computer for later manipulation or sharing. This can be video from an Internet site, or video from an external device such as a digital camera or smartphone. Video-sharing sites on the Internet make video capture cards quite popular with enterprises and Internet users alike. Video capture cards need and often come with software to aid in the processing of multimedia input. While video and sound cards are internal expansion devices, capture cards can be internal (PCIe) or external (USB).

Not all video capture cards record audio signals while processing video signals. If this feature is important to you or the end user, be sure to confirm that the card supports it. Also know that capture cards work with standard video resolutions, and specific cards might be limited on the resolutions they support. Double-check the specifications to make sure the card will meet the need, and also make sure to get reviews of the software used with the device.

Because their names are similar, it's easy to confuse the purpose of video cards and video capture cards. They do completely different things.

A video card (or graphics card) creates the visual output sent to a monitor or other display device. A video capture card takes video signals from another

source, such as a gaming console, and converts them into digital data that can be recorded and saved or streamed. Be sure to understand the difference between them!

Network Interface Card

A *network interface card (NIC)* is an expansion card that connects a computer to a network so that it can communicate with other computers on that network. It translates the data from the parallel data stream used inside the computer into the serial data stream that makes up the frames used on the network. Internal cards have a connector for the type of expansion bus on the motherboard (PCIe or PCI) and external cards typically use USB. In addition to physically installing the NIC, you need to install drivers for the NIC in order for the computer to use the adapter to access the network. Figure 2.3 shows a PCIe x1 Ethernet NIC with an RJ45 port. (Network connectors are covered in more detail in Volume 1 Chapter 6, "Networking Fundamentals .")

Many computers, especially mobile devices, have NIC circuitry integrated into their motherboards. Therefore, a computer with an integrated NIC wouldn't need to have a NIC expansion card installed unless it was faster or you were using the second NIC for load balancing (splitting the traffic between two cards to increase speed), security, or fault-tolerance (having a backup in case one fails) applications.

FIGURE 2.3 A network interface card

You will see two different types of network cards: wired and wireless. A wired card has an interface for the type of network it is connecting to (such as fiber connectors, Registered Jack 45 [RJ45] for unshielded twisted pair [UTP], antenna for wireless, or BNC for legacy coax). Wireless cards of course don't need to use wires, so they won't necessarily have one of these ports. Some do, just for compatibility or as a backup option.

Wireless NICs have the unique characteristic of requiring that you configure their connecting device before configuring the NIC. Wired NICs can generally create a link and begin operation just by being physically connected to a central hub or switch. Wireless (Wi-Fi) networks also use a central connectivity point, typically a wireless access point or wireless router, but that device must be configured with the network name and other parameters before any communication can occur. These concepts are explained in greater detail in Volume 1 Chapter 9, "Wireless and SOHO Networking." Figure 2.4 shows a PCIe wireless NIC for a desktop computer. Both antennae need to be attached for it to work—we took one off so you could see the connector. Other wireless antennas will have longer lead wires so the antenna can sit on top of the computer or a desk.

FIGURE 2.4 A wireless NIC

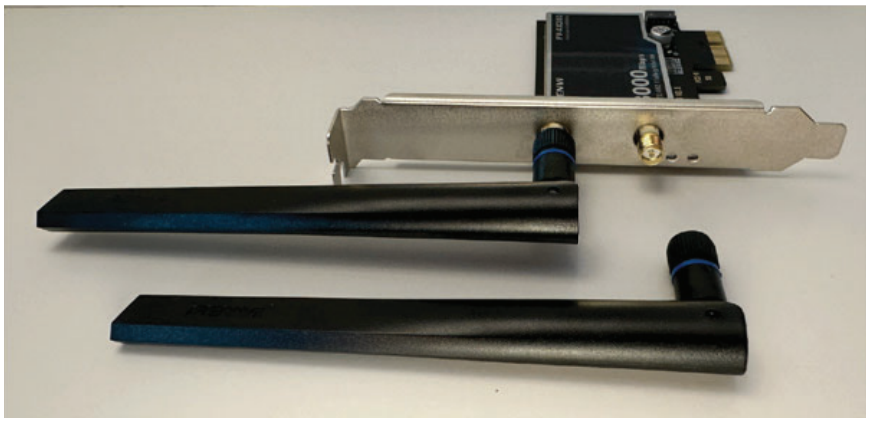

Real World Scenario

The Way of the Dinosaur. . .

Before high-speed, wireless Internet was available pretty much everywhere you went, wannabe web users needed to connect to the Internet using the plain old telephone system (POTS). Doing so required a device called a modem.

Old telephone lines were analog, and of course computers are digital. The modem got its name because it *mo*dulated and *demo*dulated analog and digital signals to enable computers to connect to each other over analog telephone lines. Modems look a bit like network cards, but the connector is smaller. Instead of the 8-pin RJ45 connector, phone lines use the 4-pin RJ11 connector. For those of you interested in history, a modem is shown in Figure 2.5. Notice it has two connectors. It's probably too hard to see in Figure 2.5, but modems with two connectors have small RJ11 connector and telephone icons etched into the faceplate. This tells you which connector to use to plug into the wall, and which one to

plug a telephone into. If the modem was in use, the telephone could not be used for voice calls. Only one connection at a time!

FIGURE 2.5 An internal analog modem

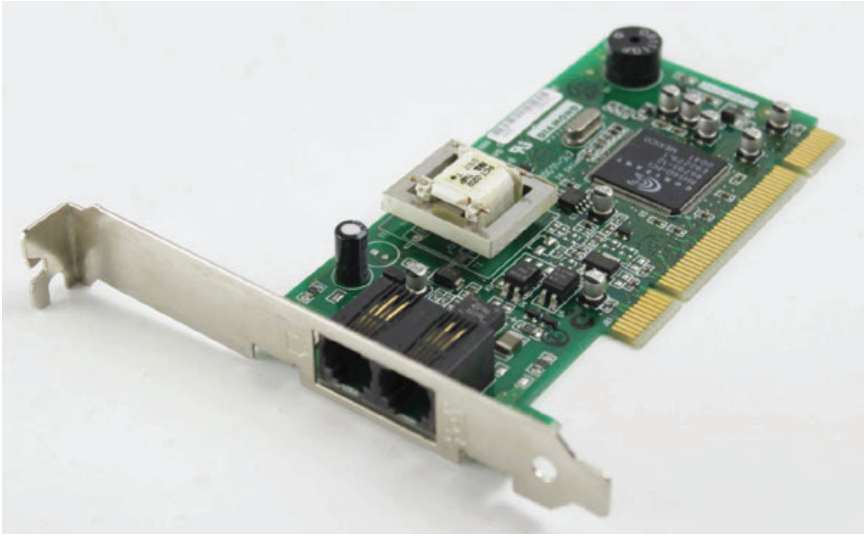

Input/Output

An *input/output card* is often used as a catchall phrase for any expansion card that enhances the system, allowing it to interface with devices that offer input to the system, output from the system, or both. The following are common examples of modern I/O cards:

- USB cards
- Storage cards, such as eSATA
- Thunderbolt cards

Figure 2.6 shows a 7-port PCIe x1 USB expansion card (left) next to an eSATA card (right). USB cards commonly come in 2-, 4-, and 7-port configurations, whereas eSATA cards often have one or two external ports. This eSATA card also has two internal SATA connectors on the top (left, in the picture) of the card. You'll also find cards that have multiple port types, such as eSATA and USB.

These cards are to be installed in a compatible slot on the motherboard. Their configuration is minimal, and it is usually completed through the operating system's Plug and Play (PnP) process. Nevertheless, check the BIOS/UEFI settings after installation for new entries in the menu structure. It's the job of the BIOS/UEFI to track all the hardware in the system and supply resources as needed. For example, a new Thunderbolt expansion card might allow you to configure whether attached Thunderbolt devices should be allowed to wake the system, how long a delay should be observed before waking the system, and various settings for how to use memory and other resources.

FIGURE 2.6 USB and eSATA expansion cards

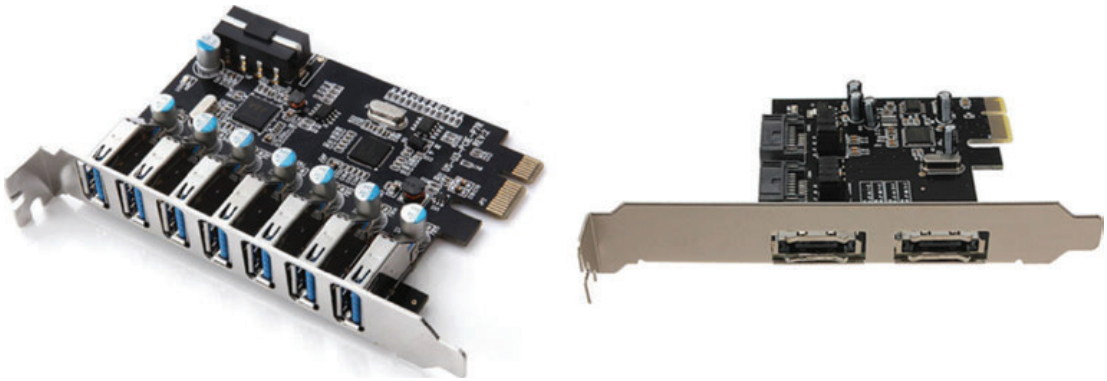

Adapter Configuration

Most adapter cards will be automatically recognized by a PnP operating system such as Windows or macOS. On the off chance that the device isn't recognized or isn't installed properly, a driver might need to be manually installed. In addition, some adapters require additional configuration of features. We will look at drivers and adapter configuration here.

Installing and Updating Device Drivers

As noted previously, when you install a new device, a PnP OS should recognize it and install a driver for you. This goes for both internal and external devices. While this process works flawlessly the vast majority of the time, it's not perfect. In addition, sometimes the default driver supplied by the OS is older or doesn't support all of the features of a device like a newer driver will. It's good practice to always have the most updated driver for every device. For some devices such as video cards, driver updates might appear every six months. For others, such as basic keyboards or mice, it might be every decade.

Drivers in Windows are updated through the Device Manager utility, shown in Figure 2.7. If a device is not functioning properly, the device icon will have a yellow exclamation point or a red X over it. If you see devices like this listed, updating the driver is a good step to resolving the problem.

Device Manager has options to see driver details, update the driver, roll back the driver (this is useful if a driver update causes problems with the device working), disable the device, and uninstall the device. To get to these options, right-click the device in the Device Manager list and click Properties. Then open the Driver tab, as shown in Figure 2.8.

Exercise 2.1 shows you how to update a device driver in Windows 11. As pre-steps for the exercise, you should do the following:

1. Go to the manufacturer's website and download the latest driver. Save it in a place that's easy to find.
2. Double-click the downloaded driver file. It will do one or both of the following things:
    a. Extract the driver files to a separate location. Again, make that location easy to find.
    b. Start installing the driver (if so, go ahead and complete the process and review Exercise 2.1 for reference).

FIGURE 2.7 Windows 11 Device Manager

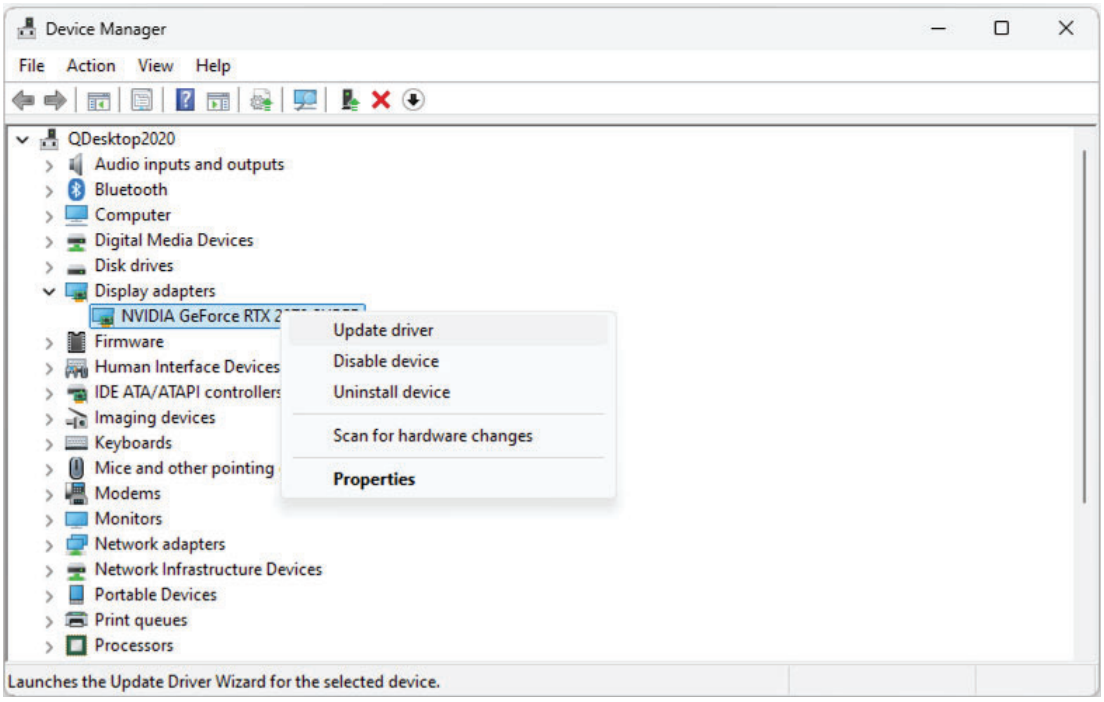

FIGURE 2.8 Video card properties

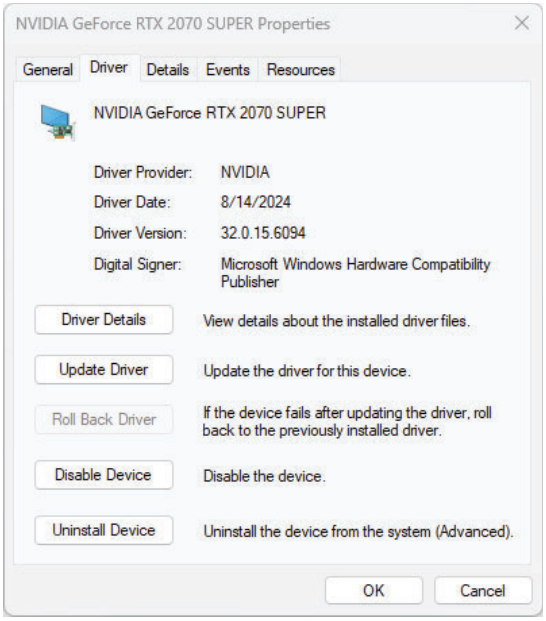
NVIDIA GeForce RTX 2070 SUPER Properties
General Driver Details Events Resources
NVIDIA GeForce RTX 2070 SUPER
Driver Provider: NVIDIA
Driver Date: 8/14/2024
Driver Version: 32.0.15.6094
Digital Signer: Microsoft Windows Hardware Compatibility Publisher
Driver Details
View details about the installed driver files.
Update Driver
Update the driver for this device.
Roll Back Driver
If the device fails after updating the driver, roll back to the previously installed driver.
Disable Device
Disable the device.
Uninstall Device
Uninstall the device from the system (Advanced).
OK
Cancel

The video driver we downloaded for Exercise 2.1 does the extraction routine and then immediately starts the installation. The installation program is shown as an example in Figure 2.9. The installation will vary by driver, but this gives you an idea of what the program will look like. Again, if you choose to install or update a driver this way, that's completely fine; you won't need to complete the steps in Exercise 2.1, though.

FIGURE 2.9 NVIDIA driver installation

NVIDIA Installer
NVIDIA Graphics Driver
Version 566.36
NVIDIA
System Check
License Agreement
Options
Install
Finish
NVIDIA Graphics Driver and NVIDIA App
NVIDIA Graphics Driver
The NVIDIA app requires data to configure your graphics card and application settings for the best gaming, content creation, and design experience. We need to know your PC's hardware and software configuration, installed games and creative/design apps, game and app settings and performance data, usage metrics of the NVIDIA app, and geographical region.
We use this data to optimize your application settings, deliver software and content updates, improve application features, redeem codes for free games, and comply with local privacy laws.
Our privacy policy explains what data we collect and what we do with it. You may configure privacy options in the NVIDIA app Settings menu.
By installing you agree to the following:
End User License Agreement
NVIDIA Privacy Policy
AGREE AND CONTINUE
CANCEL

EXERCISE 2.1

Updating a Device Driver in Windows 11

This exercise uses a video card as an example, but the process is similar for all devices. It will also work for new driver installation (provided the device is detected in Device Manager) or driver updates.

1. Open Device Manager by typing **`device manager`** into the Windows 11 search box and pressing Enter.
2. Right-click the device. You will see a menu similar to the one in Figure 2.7.
3. Click Update Driver. That will begin the driver update process shown in Figure 2.10.

EXERCISE 2.1 *(continued)*

FIGURE 2.10 Update driver in Windows

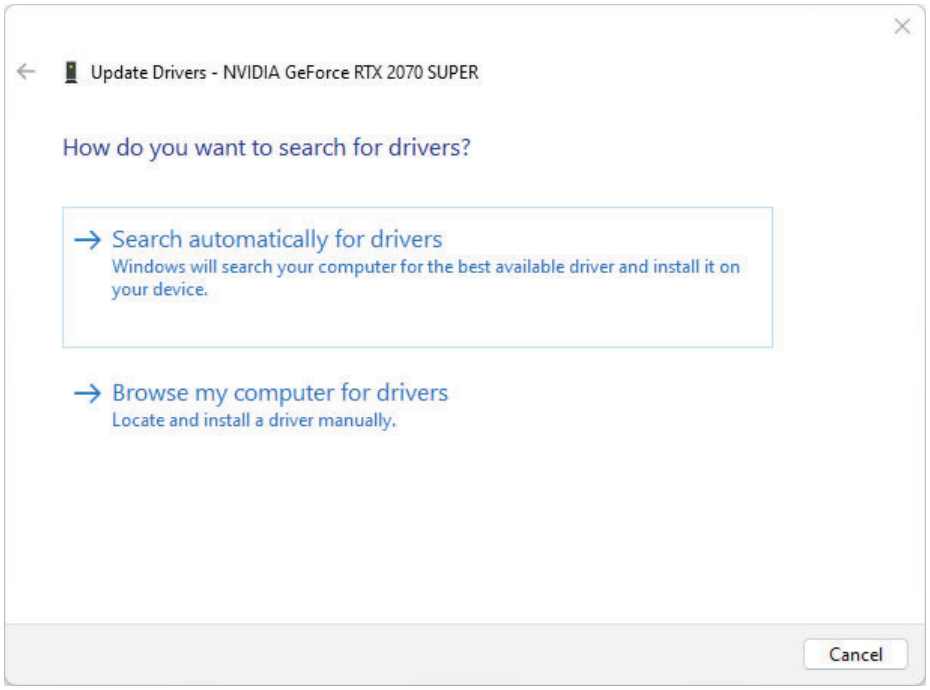
Update Drivers - NVIDIA GeForce RTX 2070 SUPER
How do you want to search for drivers?
Search automatically for drivers
Windows will search your computer for the best available driver and install it on your device.
Browse my computer for drivers
Locate and install a driver manually.
Cancel

4. Have Windows search for the driver, or specify the location where you saved the extracted driver files.
5. Follow the prompts to complete the installation.

Configuring Devices

Most devices come preconfigured with the most popular options. Some devices might not have configuration options at all, whereas others have quite a few. Some devices might have software settings configured through an app provided by the device manufacturer. If there is no app, try opening Device Manager, right-clicking the device, and choosing Properties. Again, options will vary by device. Here are some examples:

- Video cards might require you to set the resolution or enable/disable other advanced features.
- Most sound cards tend to work with no specific configuration, but advanced features will need to be implemented through the operating system or through utilities that came with the adapter.

- The functions of video capture cards are sometimes not native to the operating system and therefore come with advanced utilities that must be learned and configured before the adapters will work as expected.
- Wireless network adapters often require the installation of a screw-on antenna, which should be postponed until after the card is fully inserted and physically secured in the system. Software configuration that allows these cards to communicate with a wireless access point can be challenging for the novice. To give you an example of what a wireless NIC configuration screen may look like, take a look at Figure 2.11. We will cover some of these and others in Volume 1 Chapter 9. (Most of them you will never need to set in your entire career.)

FIGURE 2.11 Wireless NIC properties

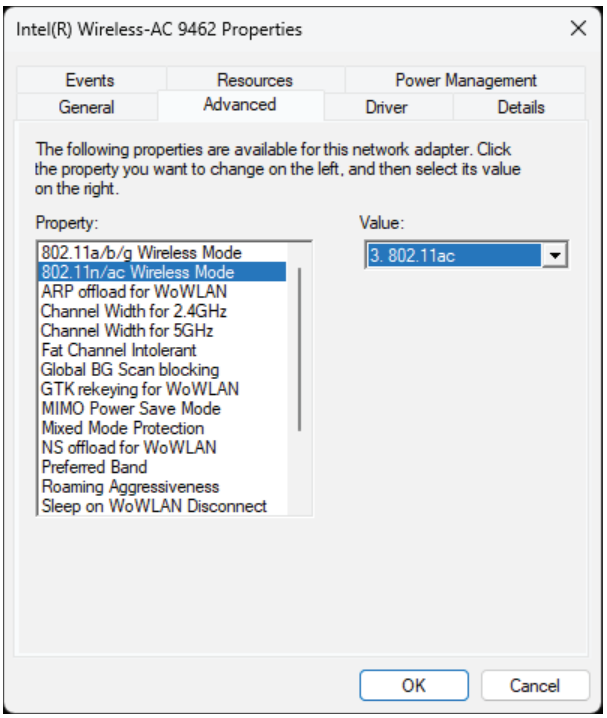

Wired network adapters tend to be easier to configure than wireless ones. Nevertheless, even wired NICs might require manual configuration of certain protocol settings, such as IP addressing, duplex, and speed, in order for them to be productive. We will cover these topics in later networking chapters. In general, installation and configuration steps for expansion cards can be summarized as follows:

1. Ensure that the computer is powered off and that you are properly grounded.
2. Install the adapter into an open slot.
3. Connect power, if needed.

   This most often applies to video cards.

4. After booting up the computer, install the driver.

   Again, Plug and Play may take care of this automatically for you.
5. If the card isn't recognized or providing the expected functionality, check the BIOS for configuration settings.
6. Check Device Manager to see if the device is recognized and if there are any error badges.
7. For other configuration options, use the utility provided by the manufacturer, if applicable.

In any event, consult the documentation provided with your adapter or the manufacturer's website for additional configuration requirements or options. The more specialized the adapter, the more likely it will come with specialty-configuration utilities.

Understanding Storage Devices

What good is a computer without a place to put everything? Storage media hold the files that the operating system needs to operate and the data that users need to save. What about saving to the cloud? The computers that make up the cloud, rather than the local computer, hold the storage media. The many different types of storage media differ in terms of their capacity (how much they can store), access time (how fast the computer can access the information), and the physical type of media used. The following types of storage devices are covered here:

- Conventional hard disk drives
- Solid-state drives
- RAID
- Removable storage, including optical drives

After that, we'll finish this section by looking at installing, removing, and configuring storage devices.

Conventional Hard Disk Drives

Hard disk drive (HDD) systems (or *hard drives* for short) are used for permanent, nonvolatile data storage and quick access. Hard drives typically reside inside the computer, where they are semi-permanently mounted with no external access (although there are external and removable hard drives) and can hold more information than other forms of storage. Conventional hard drives use a *magnetic storage* medium on rotating metal platters.

Hard drives, both conventional and solid-state (covered next), are nonvolatile storage. This means that unlike RAM or cache, they retain their data when power is shut off.

Hard drives often connect to the motherboard's SATA interface, which you learned about in Volume 1 Chapter 1, "Motherboards, Processors, and Memory." The back of the hard drive will have data and power connectors. Figure 2.12 offers a bottom view of a SATA hard drive, and shows the data and power connectors.

FIGURE 2.12 SATA hard drive connectors

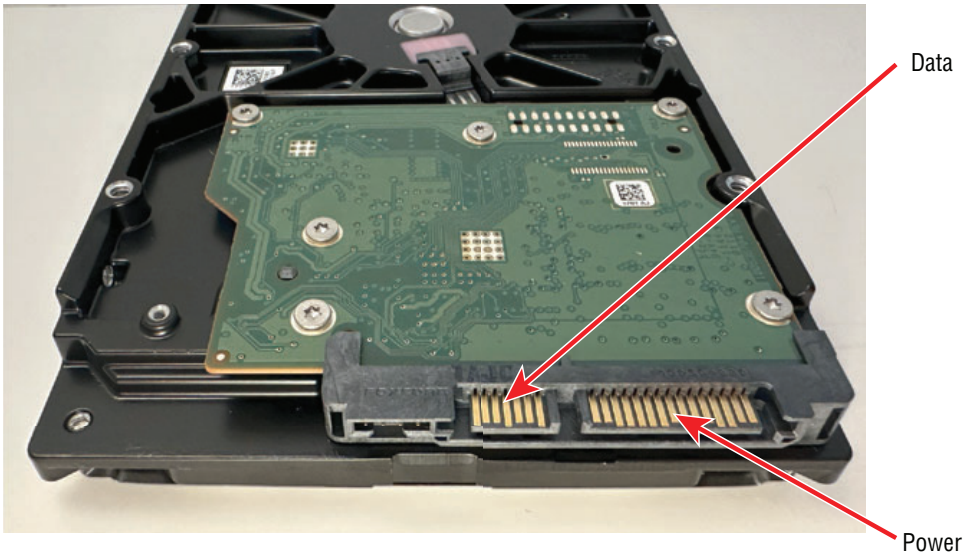

Putting Ancient Technology Aside

The predecessor to SATA for storage connections was Parallel Advanced Technology Attachment (PATA), which we briefly introduced in Volume 1 Chapter 1. PATA was used for hard drives, optical drives, and magnetic tape drives. Today, PATA hard drives are essentially obsolete because of their slow transfer speeds. Most PATA hard drives followed the ATA/100 standard, which had a maximum transfer speed of 100 MBps.

You might still run across a PATA drive, so Figure 2.13 shows you the data and power connectors for one. The data connector on the right is 40 pins and uses a flat ribbon cable. There is a missing pin, or key, in the center of the top row of pins. The 4-pin power connector on the left is called Molex (which is still an exam objective). We show the *Molex connector* from the power supply later in this chapter in Figure 2.32 and again in Volume 1

Chapter 3. In between them is the jumper block, which is used to configure the drive as primary or secondary (there can only be two PATA drives on one PATA controller).

FIGURE 2.13 PATA hard drive connectors

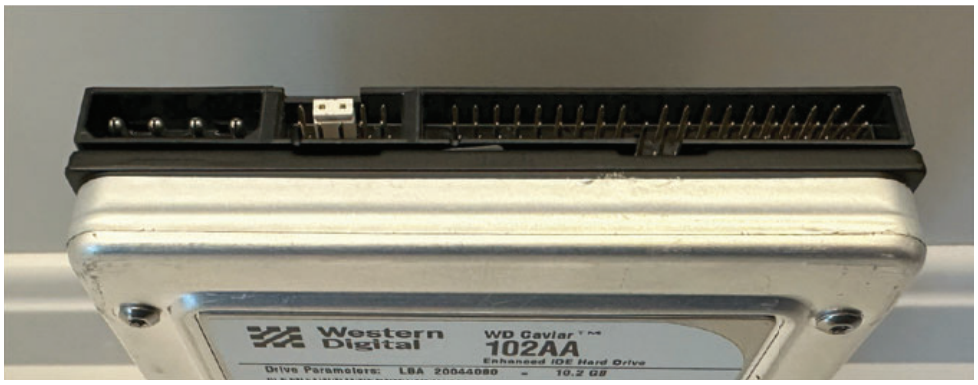

Anatomy of a Hard Drive

A hard drive is constructed in a cleanroom to avoid the introduction of contaminants into the hermetically sealed drive casing. Once the casing is sealed, most manufacturers seal one or more of the screws with a sticker warning that removal of or damage to the seal will result in voiding the drive's warranty. Even some of the smallest contaminants can damage the precision components if allowed inside the hard drive's external shell.

Inside the sealed case of the hard drive lie one or more platters, where the actual data is stored by the read/write heads. The heads are mounted on a mechanism that moves them in tandem across both surfaces of all platters. Older drives used a stepper motor to position the heads at discrete points along the surface of the platters, which spin at thousands of revolutions per minute on a spindle mounted to a hub. Newer drives use voice coils for a more analog movement, resulting in reduced data loss because the circuitry can sense where the data is located through a servo scheme, even if the data shifts due to changes in physical disc geometry. Figure 2.14 shows the internal components of a conventional hard drive.

Before a hard drive can store data, it must be prepared. The drive is separated into logical areas with a process called partitioning. After that, the partitions can be formatted with a filesystem and then it's ready to store data.

We will talk more about filesystems such as FAT32, NTFS, and others in Volume 2 Chapter 1, "Operating System Basics."

HDD Spindle Speeds

As the electronics within hard drives got faster, they were capable of requesting data at higher and higher rates. If the platters were spinning at a constant rate, however, the information could be accessed only as fast as a given fixed rate. To make information

available to the electronics more quickly, manufacturers increased the speed at which the platters spin from one generation of drives to the next, with multiple speeds coexisting in the marketplace for an unpredictable period, at least until the demand died down for one or more speeds.

FIGURE 2.14 Anatomy of a hard drive

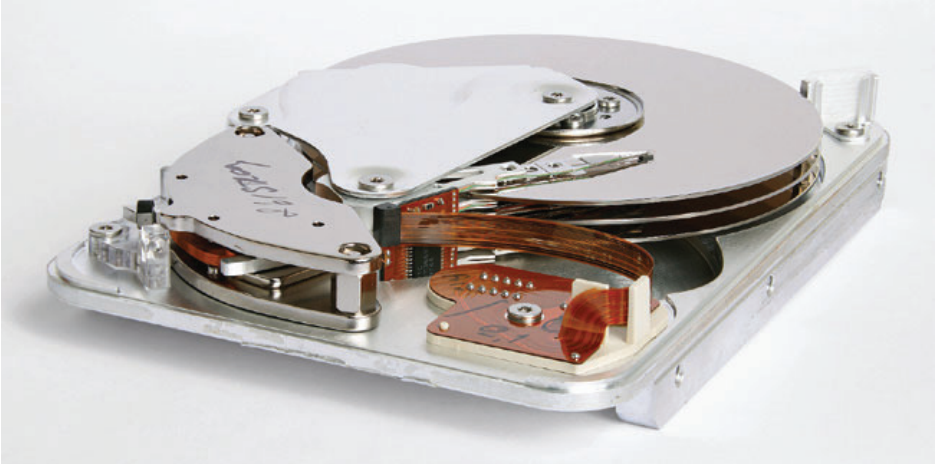

By Eric Gaba, Wikimedia Commons user Sting, CC BY-SA 3.0, `https://commons.wikimedia.org/w/index.php?curid=11278668`

Exam Objective 3.4, "Compare and contrast storage devices," calls out spindle speeds and form factors as topics to know for the exam.

The following *spindle speeds* (or spin rates) have been used in the industry for the platters in conventional magnetic hard disk drives:

- 5,400 rpm
- 7,200 rpm
- 10,000 rpm
- 12,000 rpm
- 15,000 rpm

While it is true that a higher revolutions per minute (rpm) rating results in the ability to move data more quickly, many applications do not benefit from increased disk-access speeds. As a result, you should choose only faster drives, which are also usually more expensive per byte of capacity, when you have an application for this type of performance, such as for housing the partition where the operating system resides or for very disk-intensive programs. For comparison, a 7,200 rpm SATA hard drive can sustain data read speeds of about 100 MBps, which is about the same as a PATA ATA/100 7,200 rpm drive. A 10,000 rpm (also known as 10k) SATA drive can top out around 200 MBps.

Most conventional hard drives sold are 5,400 rpm or 7,200 rpm. Faster drives were usually reserved for specialized installations.

Higher speeds also consume more energy and produce more heat. The lower speeds can be ideal in laptops, where heat production and battery usage can be issues with higher-speed drives. Even the fastest conventional hard drives are slower than solid-state drives are at transferring data.

HDD Form Factors

Physically, the most common hard drive *form factors* (sizes) are 3.5″ and 2.5″. Desktops traditionally use 3.5″ drives, whereas the 2.5″ drives are made for laptops—although most laptops today avoid using conventional HDDs. Converter kits are available to mount a 2.5″ drive into a 3.5″ desktop hard drive bay. Figure 2.15 shows the two drives together. As you can see, the 2.5″ drive is significantly smaller in all three dimensions, but it does have the same storage capacity and connectors as its bigger cousin.

FIGURE 2.15 A 3.5″ and a 2.5″ hard drive

Reading Is Like Breathing In, and Writing Is Like Breathing Out

There are two categories of conventional hard drives that you will find in the marketplace, and they differ in how they read and write data. The first is conventional magnetic recording (CMR), which is the "traditional" hard drive. Data is written in tracks (think of concentric circles on the platter), with a small gap between each track. A newer technology is called shingled magnetic recording (SMR), which overlaps tracks to increase storage capacity.

CMR drives are better for situations where data will be frequently read from and written to a drive. SMR drives are less expensive per platter and more energy efficient, but slower at transferring data. Therefore, they are best suited for things like data backups or archives.

Neither CMR nor SMR is an exam objective, but knowing details like this can help you make a better decision if you're in the market for a new drive.

Solid-State Drives

Unlike conventional hard drives, *solid-state drives (SSDs)* have no moving parts—they use the same solid-state memory technology found in the other forms of flash memory. You can think of them as big versions of the flash drives that are so common.

Because they have no moving parts, SSDs are capable of transferring data much more quickly than HDDs could ever dream of doing. Recall from the "HDD Spindle Speeds" section that a 10k SATA HDD tops out at about 200 MBps. Even the slowest SSDs will run circles around that. The true speed of an SSD will be determined, of course, by the drive itself, but also the interface to which it's attached.

And because there's no need for spinning platters and read/write heads, SSDs can be made much smaller than HDDs, making them better for laptops and portable devices. SSDs have several other advantages over their mechanical counterparts as well, including the following:

- Faster start-up and read times
- Less power consumption and heat produced
- Silent operation
- Generally more reliable because of a lack of moving parts
- Less susceptible to damage from physical shock and heat production
- Higher data density per square centimeter

The disadvantages of SSDs are as follows:

- The technology to build an SSD is more expensive per byte.
- All solid-state memory is limited to a finite number of write (including erase) operations. Lack of longevity could be an issue. This doesn't seem to be a problem with reputable drive manufacturers, but it could become a problem as a device ages.

You will find that SSDs in the market generally have lower overall capacity than HDDs. For example, it's common to find HDDs over 8 TB in size, with 24 TB drives pacing the market. Conversely, the biggest commercially available SSD (as of this writing) has been 8 TB for a few years now. As for cost, HDDs run about 2–3 cents per GB and low-end SATA SSDs are about three to four times as expensive. Faster SSDs such as NVMe drives (which we'll get to in a minute) can be from four to 10 times as expensive. Of course, prices are subject to (and it's guaranteed they will) change!

When used as a replacement for traditional HDDs, SSDs are expected to behave in a similar fashion, mainly by persistently storing data—retaining contents across a power cycle. With SSD, you can also expect to exceed the read/write speed of the HDD. SSDs can be made faster still by including a small amount of DRAM as a cache.

SSDs come in various shapes and sizes and have a few different interfaces and form factors. We will cover those in the upcoming "SSD Communication Interfaces" and "SSD Form Factors" sections.

SSDs are in Exam Objective 3.4. Be able to compare and contrast the NVMe, SATA, PCIe, and SAS communications interfaces, and the M.2 and mSATA form factors.

SSD Communication Interfaces

The advent of the SSD was a major advancement for the computer industry. Solid-state drives are basically made from the same circuitry that RAM is, and they are really, really fast. When they were first on the market, the limitation in the system was the SATA controller that most hard drives were plugged into. So as enterprising computer engineers are known to do, some started looking into ways to overcome this barrier. The result is that there have been more interfaces designed for storage devices, with much faster speeds. We will look at four of them here: SATA, PCIe, NVMe, and SAS.

SATA

At this point, SATA is a technology you should be somewhat familiar with. We covered it in Volume 1 Chapter 1, and the hard drive interface is shown earlier in this chapter in Figure 2.12. SATA is a bit unique among the other subjects in this SSD section because it can support mechanical hard drives as well as SSDs.

SSDs came onto the market in the mid-1990s before SATA was even a thing, but they were limited in popularity because of the major bottleneck—the PATA communications interface. Once SATA came along in the early 2000s, the two technologies felt like glorious companions. SATA 1.x could transfer data at 150 MBps, which was a lot faster than the conventional hard drives at the time. (The most common conventional standard was ATA/100, which maxed out at 100 MBps.) Then, of course, came SATA 2.x (300 MBps) and eventually SATA 3.x (600 MBps). Comparatively, conventional hard drives appear to be painfully slow. And they are. Table 2.1 summarizes the standards and speeds.

TABLE 2.1 SATA standards and speeds

Version	Nickname	Date	Speed
1.x	SATA-150	2003	150 MBps
2.x	SATA-300	2004	300 MBps
3.x	SATA-600	2009	600 MBps

Keep in mind that of all the SSD technologies we discuss in this chapter, the SATA interface is the slowest of them. So while SATA SSDs are about 6 times faster than conventional HDDs, there is a lot of performance upside. SATA SSDs are still popular today because they are plentiful and cheap (compared to other SSDs), and motherboards usually have more SATA connectors than any other type of hard drive connector.

PCIe

Peripheral Component Interconnect Express (PCIe) is another technology that was covered in Volume 1 Chapter 1 and is used for SSDs. PCIe was first introduced in 2002, technically a year before SATA, and both technologies took a little while to get widely adopted. Like SATA, PCIe has gone through several revisions, with each version being faster than the previous one. Table 2.2 shows the throughput of PCIe versions.

TABLE 2.2 PCIe standards and transfer rates

Version	Date	Transfer rate	Throughput per lane (one direction)	Total x16 throughput (bidirectional)
1.0	2003	2.5 GT/s	250 MBps	8 GBps
2.0	2006	5.0 GT/s	500 MBps	16 GBps
3.0	2010	8.0 GT/s	1 GBps	32 GBps
4.0	2017	16.0 GT/s	2 GBps	64 GBps
5.0	2019	32.0 GT/s	4 GBps	128 GBps
6.0	2022	64.0 GT/s	8 GBps	256 GBps

PCIe 6.0 was released in 2022, and PCIe 7.0 is expected sometime in late 2025. As you might expect, it will double the speed of PCIe 6.0 and be compatible with older PCIe standards. Devices supporting the newest PCIe specifications typically take two to three years to hit the market following a new version release.

Looking at Table 2.2, something might immediately jump out at you. The transfer rate is specified in gigatransfers per second (GT/s). This unit of measure isn't very commonly used, and most people talk about PCIe in terms of its data throughput. Before moving on, though, let's take a quick look at what GT/s means. (It's highly unlikely you will be tested on this, but you might be curious.)

PCIe is a serial bus that embeds clock data into the data stream to help the sender and receiver keep track of the order of transmissions. Two clock bits are used for every 8 bits of data. Therefore, PCIe uses what's called "8b/10b" encoding—8 bits of data are sent in a

10-bit bundle, which is decoded at the receiving end. Gigatransfers per second refers to the total number of bits sent, whereas the throughput data you might be more used to seeing refers to data only and not the clock.

A little math using PCIe 1.0 might help illustrate the transfer rate versus data throughput. PCIe 1.0 has a transfer rate of 2.5 billion (giga) bits per second. Because it uses 8b/10b encoding, you then multiply by 0.8 (8/10) to get to bits of data, which is 2 trillion. But that's bits. Divide that by 8 to turn it into bytes, and you are left with 250 million bytes, or 250 MBps. Again, don't expect to see this on the test, but we know that some of you like these technical details!

Also remember that PCIe cards can have different numbers of channels—x1, x2, x4, x8, and x16. So, for example, a PCIe 3.0 x4 card will have total data throughput of 8 GBps (1 GBps for one lane in one direction, times two for bidirectional, times four for the four channels).

What does all of this mean for SSDs? First, take a look at a picture of a PCIe SSD in Figure 2.16. This is an Intel SSD 750, and here are some of its specs:

- **Interface:** PCIe 3.0 x4 – NVMe
- **Capacity:** 1.2 TB
- **Sequential read:** 2.4 GBps
- **Sequential write:** 1.2 GBps
- **Endurance:** 70 GB of writes per day for five years

FIGURE 2.16 Intel PCIe SSD

This drive is a few years old, but it's perfect for our discussion. Based on Table 2.2, PCIe 3.0 x4 has a maximum throughput of 8 GBps. This drive doesn't get close to that level, but it still reads about four times faster than the best SATA drives. This is an example of the difference between theoretical maximums of standards versus practical realities of creating hardware. Also note its rated endurance. Recall that SSDs have a finite lifespan, although it will take your average user quite some time to hit this mark.

Remember that, unlike SATA, which is designed for storage devices such as hard drives and optical drives, PCIe is more of a universal connector. Video cards, sound cards, network cards, and many other devices use PCIe slots as well. Before attempting to install a PCIe hard drive into a computer, make sure that there's an open PCIe slot of the appropriate type first!

Also remember that PCIe cards can get power from the slot itself. Devices that require a lot of power such as video cards will have their own power connector from the power supply, but PCIe hard drives don't need anything other than what the PCIe slot provides.

NVMe

Created by a consortium of manufacturers, including Intel, Samsung, Dell, SanDisk, and Seagate, and released in 2011, *Non-Volatile Memory Express (NVMe)* is an open standard designed to optimize the speed of data transfers. Unlike SATA and PCIe, NVMe isn't related to a specific type of connector. Said another way, there is no NVMe connector—think of it as a nonvolatile memory chip that can be used in SATA, PCIe, or M.2 (which we will cover in the next section) slots. Figure 2.17 shows a 1 TB Western Digital M.2 NVMe SSD.

FIGURE 2.17 M.2 NVMe SSD

NVMe drives are incredibly fast—current NVMe SSDs can support data reads of up to 3.5 GBps, provided that the interface they are plugged into supports it as well, of course. An NVMe SATA 3 SSD will still be limited to 600 MBps; older PCIe versions in theory might have limitations, but you're not going to find any PCIe 1.0 or 2.0 NVMe SSDs, so it's not a problem.

One potential issue you might see with NVMe hard drives is that in order to be a boot drive, the motherboard must support it. If the motherboard has a built-in M.2 slot, odds are that the BIOS will support booting from an NVMe drive. If you are adding it using a PCIe slot, the BIOS might not be able to boot from that drive. Always check the motherboard documentation to ensure it supports what you're trying to do.

SAS

When the Small Computer System Interface (SCSI) standard came out in 1986, it was among the fastest storage transfer technologies available, delivering 5 MBps throughput. In the early-to-mid 1990s, it was found mostly in servers. If someone ran SCSI in their home system, you knew that they were serious about high performance. Over time, revisions boosted the speed to 320 MBps (a later implementation could hit 640 MBps, but it was not widely adopted). Faster speeds were still needed though, and thanks to advancements in USB, serial technologies were showing more promise for long-term success.

Switching over to a serial standard, SCSI developers released *Serial Attached SCSI (SAS)* in 2003. It featured a solid 3 Gbps bandwidth, which is 375 MBps (see the "Converting the Numbers" sidebar in this section). At the time, the primary competing storage standards were PATA (commonly 100 MBps, with some 133 and 167 MBps implementations) and SATA 1.x (150 MBps). So SAS was very fast at the time. Subsequent versions doubled the speed (where have you heard that before?); Table 2.3 shows versions, release dates, and throughput. SAS-5 is supposedly under development, offering 45 Gbps throughput, but there is no official release date set.

TABLE 2.3 SAS standards

Version	Released	Throughput
SAS-1	2003	3 Gbps (375 MBps)
SAS-2	2009	6 Gbps (750 MBps)
SAS-3	2013	12 Gbps (1500 MBps or 1.5 GBps)
SAS-4	2017	24 Gbps (3000 MBps or 3.0 GBps)

Converting the Numbers

Notice that SAS speeds are quoted in gigabits per second (Gbps), whereas a lot of other standards are quoted in gigabytes per second (GBps) or even megabytes per second (MBps). We don't throw around these terms to confuse you—rather, we use what the standards' developers use. And yes, SCSI is in MBps and SAS is Gbps even though they were developed by the same group.

Here are some conversions, for reference:

- 125 MBps equals 1 Gbps.
- 8 Gbps equals 1 GBps.
- 1000 MBps equals 1 GBps.

Today, SAS is found mostly in servers and dedicated storage devices. You can purchase PCIe SAS expansion cards for desktop PCs, but with M.2 drives commonly available, there's not much point.

SAS hard drives come in 2.5″ and 3.5″ form factors, and here's the sometimes tricky part—the connector on the hard drive is nearly identical to a SATA connector. The only difference is that a SAS drive has a plastic block in between the data and power connectors, whereas SATA has a gap between them. SAS and SATA are also electronically compatible. SATA drives, which are less expensive than SAS drives, can be connected to a SAS controller and usually work properly, albeit at slower SATA speeds. SAS drives, however, cannot be connected to a SATA controller. Additionally, you should never mix SAS and SATA drives in the same hard drive array (logical grouping of hard drives) lest you want to deal with performance inconsistencies.

SSD Form Factors

Whereas a *communications interface* is the method the device uses to communicate with other components, a *form factor* describes the shape and size of a device. The two SSD form factors you need to know for the A+ exam are mSATA and M.2.

mSATA

The Serial ATA International Organization has developed several specifications—you've already been introduced to SATA and eSATA. Next on the list is a form factor specifically designed for portable devices such as laptops and smaller—*mini-Serial ATA (mSATA)*. mSATA was announced in 2009 as part of SATA version 3.1 and hit the market in 2010.

mSATA uses the same physical layout as the Mini PCI Express (mPCIe) standard, and both have a 30mm-wide 52-pin connector. The wiring and communications interfaces between the two standards are different though. mSATA uses SATA technology, whereas mPCIe uses PCIe. In addition, mPCIe card types are as varied as their larger PCIe cousins are, including video, network, cellular, and other devices. mSATA, on the other hand, is dedicated to storage devices based on SATA bus standards. mSATA cards come in 30mm × 50.95mm full-size and 30mm × 26.8mm half-size cards. Figure 2.18 shows a full-sized mSATA SSD on top of a 2.5″ SATA SSD.

The wiring differences between mSATA and mPCIe can pose some interesting challenges. Both types of cards will fit into the same slot, but it depends on the motherboard as to which is supported. You might have heard us say this before, but as always, check the motherboard's documentation to be sure.

FIGURE 2.18 mSATA SSD and a 2.5″ SATA SSD

Source: Vladsinger / Wikimedia Commons / CC BY SA 3.0

M.2

Originally developed under the name Next Generation Form Factor (NGFF), *M.2* (pronounced "M dot 2") was born from the desire to standardize small form factor SSD hard drives. We touched briefly on M.2 in Volume 1 Chapter 1, and we mentioned that although M.2 is primarily used for hard drives, it supports other types of cards such as Wi-Fi, Bluetooth, Global Positioning System (GPS), and near-field communication (NFC) connectivity, as well as PCIe and SATA connections. It's a form factor designed to replace the mSATA standard for ultra-small expansion components in laptops and smaller devices. Whereas mSATA uses a 30mm 52-pin connector, M.2 uses a 66-pin connector, most commonly 22mm wide.

One interesting connectivity feature of M.2 is that the slots and cards are keyed such that only a specific type of card can fit into a certain slot. The keys are given letter names to distinguish them from each other, starting with the letter *A* and moving up the alphabet as the location of the key moves across the expansion card. Table 2.4 explains the slot names, some interface types supported, and common uses.

TABLE 2.4 M.2 keying characteristics

Key	Common interfaces	Uses
A	PCIe x2, USB 2.0	Wi-Fi, Bluetooth, and cellular cards
B	PCIe x2, SATA, USB 2.0, USB 3.0, audio	SATA and PCIe x2 SSDs

Key	Common interfaces	Uses
E	PCIe x2, USB 2.0	Wi-Fi, Bluetooth, and cellular cards
M	PCIe x4, SATA	PCIe x4 SSDs

Another interesting feature of the cards is that they are also named based on their size. For example, you will see card designations such as 1630, 2242, 2280, 22110, or 3042. The first two numbers refer to the width, and the rest to the length (in millimeters) of the card.

Let's look at some examples. Figure 2.19 shows four different M.2 cards. From left to right, they are as follows:

- A- and E-keyed 1630 Wi-Fi and Bluetooth adapter
- M-keyed 2230 SSD
- B- and M-keyed 2242 SSD
- M-keyed 2280 SSD

Of the four, only an M-keyed SSD can get the fastest speeds, because it supports PCIe x4. SSDs on the market are keyed B, M, or B+M. A B-keyed or M-keyed SSD won't fit in a B+M socket. A B+M keyed drive will fit into a B socket or an M socket, however.

FIGURE 2.19 Four M.2 cards

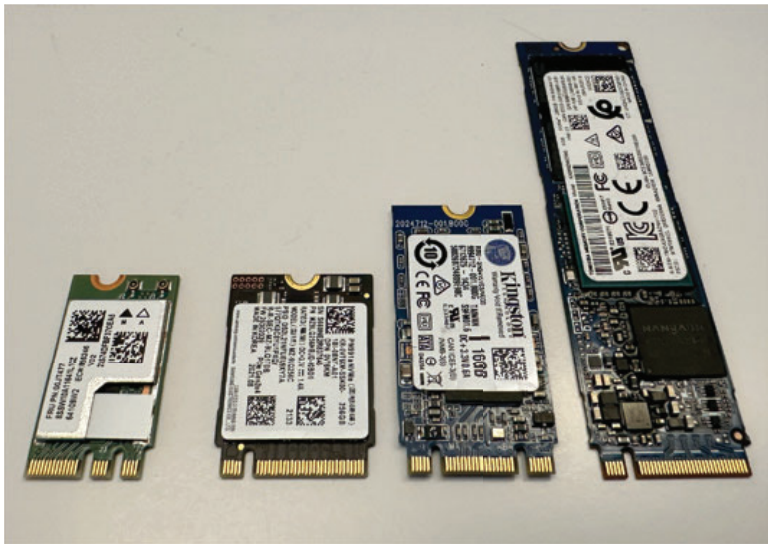

Maximum speeds are always increasing. For reference, at the time of writing, some of the fastest 4 TB NVMe M.2 SSDs support PCIe 5.0 x4 and have read speeds up to 14.1 GBps and write speeds up to 12.6 GBps.

Figure 2.20 shows a motherboard with two M.2 slots. The bottom one is hidden by its long, narrow metal cover. These function as protection and heat spreaders and get replaced over a card when it's added to the system. Be sure to take any plastic cover off the bottom of the spreader before putting it back over an M.2 card! The top one is an uncovered M-keyed slot. We put a highlighted box around the numbered post holes for 42mm, 60mm, 80mm, and 110mm cards. Use only the post hole needed for the length of card you're installing (in Figure 2.20, the post is screwed into the 80mm hole), and then the 110 post for the heat spreader. Some high-end SSDs come with their own heatsinks—in those cases, it's not necessary to reinstall the original cover. Like PCIe, M.2 devices derive power from the motherboard connector.

FIGURE 2.20 M.2 slots

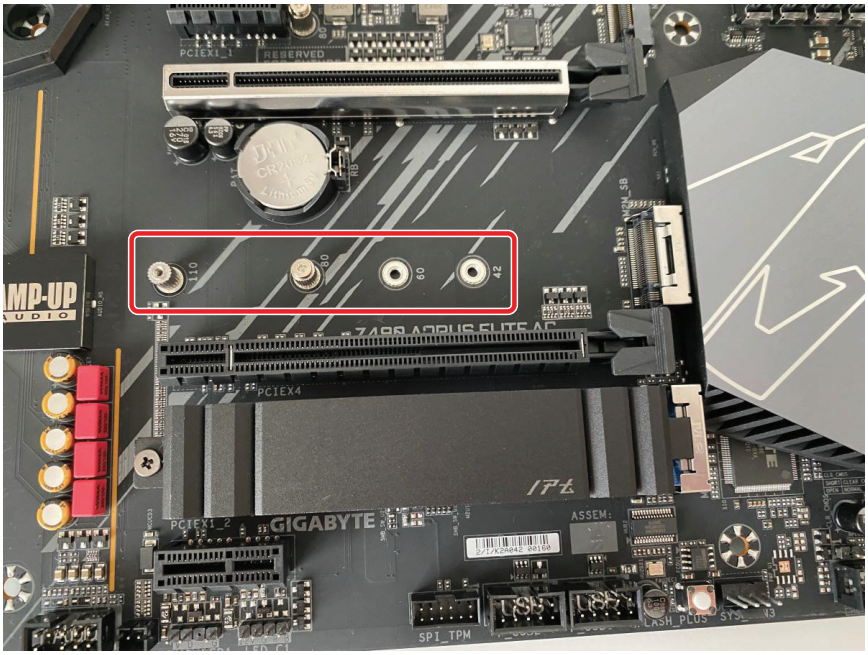

As mentioned earlier, M.2 is a form factor, not a bus standard. M.2 supports SATA, USB, and PCIe buses. What does that mean for M.2 hard drives? It means that if you purchase an M.2 SATA hard drive, it will have the same speed limitation as SATA III, or about 600 MBps. That's not terrible, but it means that the primary advantage of an M.2 SATA drive versus a conventional SATA SSD is size. An M.2 PCIe hard drive is an entirely different story. PCIe, you will recall, is much faster than SATA. A PCIe 2.0 x1 bus supports one-way data transfers of 500 MBps. That is close to SATA III speed, and it's only a single lane for an older standard. NVMe M.2 drives kick up the speed even further. If you want the gold standard for hard drive speed, NVMe M.2 is the way to go.

RAID

Multiple hard drives can work together as one system, or array, providing increased performance (faster disk reads and writes), fault tolerance (protection against one or more disks failing), or both. Such systems are called *Redundant Array of Independent (or Inexpensive) Disks (RAID)*. RAID can be implemented in software, such as through the operating system, or in hardware, such as through the motherboard BIOS or a RAID hardware controller. Hardware RAID is more efficient and offers higher performance but at an increased cost.

Exam Objective 3.4 mentions RAID. You need to be able to compare and contrast RAID 0, 1, 5, 6, and 10.

There are several types of RAID. The following are the most commonly used RAID levels:

RAID 0 *RAID 0* is also known as *disk striping*, where a striped set of equal space from at least two drives creates a larger volume. This is in contrast to unequal space on multiple disks being used to create a simple *volume set*, which is not RAID 0. RAID 0 doesn't provide the fault tolerance implied by the *redundant* component of the name. Data is written across multiple drives, so one drive can be reading or writing while another drive is doing the same. This makes for faster data access. If any one of the drives fails, however, all content is lost. Some form of redundancy or fault tolerance should be used in concert with RAID 0.

RAID 1 Also known as *disk mirroring*, *RAID 1* is a method of producing fault tolerance by writing all data simultaneously to two separate drives. If one drive fails, the other contains all of the data, and it will become the primary drive. Disk mirroring doesn't help access speed, however, and the cost is double that of a single drive. If a separate host adapter is used for the second drive, the term *duplexing* is attributed to RAID 1. Only two drives can be used in a RAID 1 array.

RAID 5 *RAID 5* combines the benefits of both RAID 0 and RAID 1, creating a redundant striped volume set. Sometimes you will hear it called a *stripe set with parity*. Unlike RAID 1, however, RAID 5 does not employ mirroring for redundancy. Each stripe places data on *n*-1 disks, and parity computed from the data is placed on the remaining disk. The parity is interleaved across all the drives in the array so that neighboring stripes have parity on different disks. If one drive fails, the parity information for the stripes that lost data can be used with the remaining data from the working drives to derive what was on the failed drive and to rebuild the set once the drive is replaced.

The same process is used to continue to serve client requests until the drive can be replaced. This process can result in a noticeable performance decrease, one that is predictable because all drives contain the same amount of data and parity. Furthermore,

the loss of an additional drive results in a catastrophic loss of all data in the array. Note that while live requests are served before the array is rebuilt, nothing needs to be computed for stripes that lost their parity. Recomputing parity for these stripes is required only when rebuilding the array. A minimum of three drives is required for RAID 5. The equivalent space of one drive is lost for redundancy. The more drives in the array, the less of a percentage this single disk represents. Because of the parity stripe, the total capacity of a RAID 5 array is the sum of the capacities of all disks in the array, minus one disk. If more than one drive fails, the RAID 5 array will fail as well.

Figure 2.21 illustrates RAID 1 and RAID 5.

FIGURE 2.21 RAID 1 and RAID 5

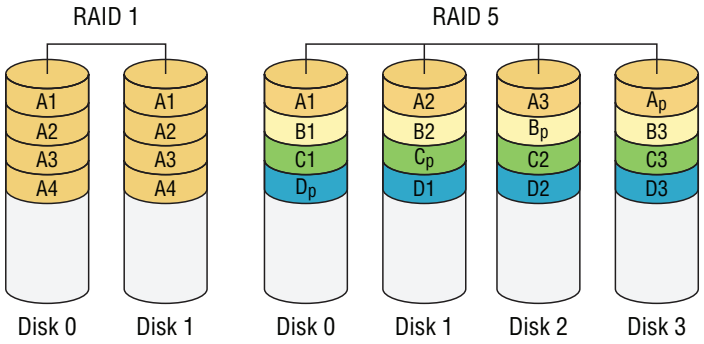

RAID 6 *RAID 6*, also called *double-parity RAID*, is similar to RAID 5. The biggest difference is that RAID 6 uses a second parity stripe, meaning that it can recover from multiple disk failures. Because of this, it's considered more fault-tolerant than RAID 5 and is a better choice for mission-critical data. It will be a little slower than RAID 5, though, because two parity stripes need to be calculated. A minimum of four hard disks is required to implement RAID 6. The total storage capacity of the array will be the sum of all drive capacities minus two drives, to account for the double-parity stripes.

RAID 10 Also known as *RAID 1+0, RAID 10* adds fault tolerance to RAID 0 through the RAID 1 mirroring of each disk in the RAID 0 striped set. Its inverse, known as *RAID 0+1*, mirrors a complete striped set to another striped set just like it. Both of these implementations require a minimum of four drives and, because of the RAID 1 component, use half of your purchased storage space for mirroring.

Table 2.5 summarizes the key data for RAID systems.

TABLE 2.5 RAID characteristics

Version	Nickname	Fault tolerant?	Min. drives	Can recover from X drive failures	Capacity
RAID 0	Disk striping	No	2	0	Add total space from all drives.

Version	Nickname	Fault tolerant?	Min. drives	Can recover from X drive failures	Capacity
RAID 1	Disk mirroring	Yes	2	1	Add total space from all drives, divide in half.
RAID 5	Stripe set with parity	Yes	3	1	All drives minus one drive.
RAID 6	Double parity	Yes	4	2	All drives minus two drives.
RAID 10	RAID 1+0	Yes	4	2*	Add total space from all drives, divide in half.

* The number of drive failures a RAID 10 system can sustain depends on a few factors. If both mirrored or both original drives fail, the data can be recovered. But if Disk 1 of both the original set and the mirrored set fail, the data is not recoverable. Always have a data backup!

Removable Storage and Media

Thus far we've focused on storage media that is internal to a PC, but external and removable storage options exist as well. Among the other types of storage available are flash drives, memory cards, optical drives, and external hard drives. The following sections present the details about removable storage solutions.

You need to know about removable storage such as flash drives and memory cards, as well as optical drives, as part of Exam Objective 3.4.

Flash Memory

Once used only for primary memory, the same components that sit on your motherboard as RAM can be found in various physical sizes and quantities among today's solid-state storage solutions. These include older removable and nonremovable flash memory mechanisms, Secure Digital (SD) and other memory cards, and USB flash drives. Each of these technologies has the potential to store reliably a staggering amount of information in a minute form factor. Manufacturers are using innovative packaging for some of these products to provide convenient transport options (such as keychain attachments) to users. Additionally, recall the SSD alternatives to magnetic hard drives mentioned earlier in this chapter.

For many years, modules known as flash memory have offered low- to mid-capacity storage for devices. The name comes from the concept of easily being able to use electricity to alter the contents of the memory instantly. The original flash memory is still used in devices that require a nonvolatile means of storing critical data and code often used in booting the device, such as routers and switches.

For example, Cisco Systems uses flash memory in various forms to store its Internetwork Operating System (IOS), which is accessed from flash during bootup and, in certain cases, throughout operation uptime and therefore during an administrator's configuration sessions. Lesser models store the IOS in compressed form in the flash memory device and then decompress the IOS into RAM, where it is used during configuration and operation. In this case, the flash memory is not accessed again after the boot-up process is complete, unless its contents are being changed, as in an IOS upgrade.

The following sections explain a bit more about today's most popular forms of flash memory: USB flash drives and memory cards.

USB Flash Drives

USB *flash drives* are incredibly versatile and convenient devices that enable you to store large quantities of information in a very small form factor. Many such devices are merely extensions of the host's USB connector, extending out from the interface but adding little to its width, making them easy to transport, whether in a pocket or a laptop bag. Figure 2.22 shows a few flash drives.

FIGURE 2.22 USB flash drives

USB flash drives capitalize on the versatility of the USB interface, taking advantage of Windows' Plug and Play, AutoPlay, and Safely Remove Hardware features and the physical connector strength. Upon insertion, these devices announce themselves to Windows File Explorer as removable drives, and they show up in the Explorer window with a drive letter. This software interface allows for drag-and-drop copying and most of the other Explorer functions performed on standard drives.

Flash drives exist in a variety of capacities and price points. You can find 32 GB to 64 GB drives for under $5, while drives in the 512 GB to 1 TB range might cost $50 or more. USB flash drives may be Type-A (the old rectangular connecter) or Type-C (the newer oval-shaped one).

SD and Other Memory Cards

Today's smaller devices require some form of removable solid-state memory that can be used for temporary and permanent storage of digital information. Modern electronics, as well as most contemporary digital still cameras, use some form of removable *memory card* to store still images permanently or until they can be copied off or printed out. Of these, the *Secure Digital (SD)* format has emerged as the preeminent leader of the pack, which includes the older MultiMediaCard (MMC) format on which SD is based. Both of these cards measure 32mm × 24mm, and slots that receive them are often marked for both. The SD card is slightly thicker than the MMC and has a write-protect notch (and often a switch to open and close the notch), unlike MMC.

Even smaller devices, such as mobile phones, have an SD solution for them. One of these products, known as *miniSD*, is slightly thinner than SD and measures 21.5mm × 20mm. The other, *microSD*, is thinner yet and only 15mm × 11mm. Both of these reduced formats have adapters that allow them to be used in standard SD slots. Figure 2.23 shows an SD card and a microSD card next to a ruler based on inches. There are other formats as well, such as CompactFlash (CF), which is used in digital cameras, but they are far less common than SD cards.

FIGURE 2.23 Typical SD cards

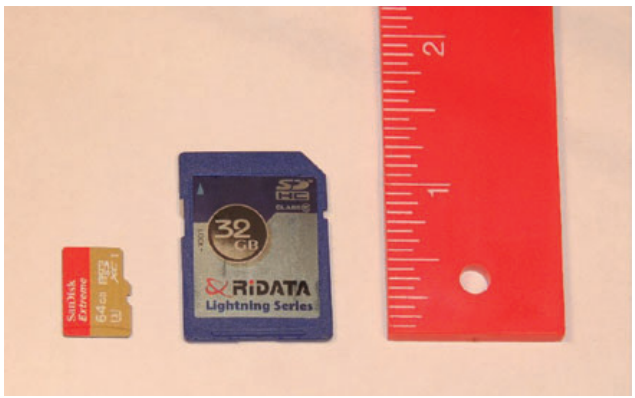

Some printers and laptops have built-in memory card readers. If not, users can purchase a USB-attached card reader like the one shown in Figure 2.24. The one in Figure 2.24 is simple, with slots for SD and microSD cards, as well as a USB-A port for flash drives.

Hot-Swappable Devices

Many of the removable storage devices mentioned are *hot-swappable*. This means that you can insert and remove the device with the system powered on. Most USB-attached devices without a filesystem fall into this category. Non–hot-swappable devices, in contrast, either cannot have the system's power applied when they are inserted or removed or have some sort of additional conditions for their insertion or removal. One subset is occasionally

referred to as cold-swappable, the other as warm-swappable. The system power must be off before you can insert or remove cold-swappable devices. An example of a cold-swappable device is anything connected to a SATA connector on the motherboard.

FIGURE 2.24 A USB-attached SD card reader

Warm-swappable devices include USB flash drives and external drives that have a filesystem. Windows and other operating systems tend to leave files open while accessing them and write cached changes to them at a later time, based on the algorithm in use by the software. Removing such a device without using the Safely Remove Hardware and Eject Media utility can result in data loss. However, after stopping the device with the utility, you can remove it without powering down the system—hence, the *warm* component of the category's name. These are officially hot-swappable devices.

Hardware-based RAID systems benefit from devices and bays with a single connector that contains both power and data connections instead of two separate connectors. This is known as *Single Connector Attachment (SCA)*. SCA interfaces have ground leads that are longer than the power leads so that they make contact first and lose contact last. SATA power connectors are designed in a similar fashion for the same purpose. This arrangement ensures that no power leads make contact without their singular ground leads, which would often result in damage to the drive. Drives based on SCA are hot-swappable. RAID systems that have to be taken offline before drives are changed out, but the system power can remain on, are examples of warm-swappable systems.

Optical Drives

The final category of storage devices we will look at is optical drives. They get their name because instead of storing data using magnetic fields like conventional HDDs, they read and write data with the use of a laser. The laser scans the surface of a spinning plastic disc, with data encoded as small pits and bumps in the track of the disc.

Optical storage devices began earnestly replacing floppy diskette drives in the late 1990s. While they had far higher capacity than floppies, they were never intended to replace HDDs, because HDDs greatly exceed the capacity and read/write performance of optical drives. For about 20 years, optical drives were practically required components, but today they are far less common. The popularity of high-speed Internet access, streaming services, and

high-capacity USB drives that don't require a separate hardware reader have relegated them to low-priority peripherals.

The most advanced optical storage technology used is the *Blu-ray Disc (BD)* drive. It replaced the *digital versatile disc (DVD)*, also called *digital video disc drive*, which in turn replaced the *compact disc (CD)* drive. Each type of optical drive can also be expected to support discs of the technologies that came before it. Figure 2.25 shows an example of an internal 5¼″ DVD-ROM drive, which also accepts CD-ROM discs. Modern optical drives are indistinguishable from older ones, aside from obvious markings concerning their capabilities. External drives that connect via USB are more popular (and portable!) than their internal cousins.

FIGURE 2.25 A DVD-ROM drive

The original compact disc was intended to be read-only. Data was burned, or written, to the disc once and that was it. Based on consumer demand, manufacturers started creating re-recordable discs. The basic CD-ROM that could both read and write was called a *burner*, a nickname that stuck with future generations of optical writers as well.

There's a lot of interesting history behind the development of each iteration of optical discs, complete with competing standards, intrigue, and infighting. Because optical discs are becoming more and more obscure, we'll cover just some of the key points to know:

- CDs could only be written to one side. Eventually that became known as single-sided (SS).
- DVDs used better technology that allowed the ability to write on both sides (dual-sided [DS]), and have two layers of data on one side (dual-layer [DL]).
- Dual-sided DVDs doubled the capacity of single-sided ones. But in dual-layer DVDs, accommodations needed to be made to avoid overlapping tracks, resulting in about a 10 percent reduction in capacity.
- Further technology enhancements, such as using a violet laser instead of a red one and improved optics, allowed BDs to have much greater capacity. (The name Blu-ray is a bit of creative license from the change in laser color.)

The last piece to cover is transfer rates. Each standard has its own base data transfer speed, which is its 1x speed. Devices in the market will read and write at several multiples of the base speed, which we call the common transfer rate—it reflects the fastest devices typically available. Also note that when burning discs, the transfer rate will be much slower than what's listed here. For example, you might find a BD burner that reads at 16x but burns at 6x. Table 2.6 outlines some basic specifications of the different optical technologies.

TABLE 2.6 Optical disc and drive specifications

Technology	Format	Standard capacity	1x transfer rate	Common transfer rate
CD	SS	700 MB	150 Kbps	7.8 Mbps (52x)
DVD	SS, SL	4.7 GB	1.4 Mbps	33.6 Mbps (24x)
DVD	SS, DL	8.5 GB		
DVD	DS, SL	9.4 GB		
BD	SS, SL	25 GB	4.5 Mbps	72 Mbps (16x)
BD	SS, DL	50 GB		
BD	DS, DL	100 GB		

SS = single-sided; DS = double-sided; SL = single-layer; DL = dual-layer

Installing, Removing, and Configuring Storage Devices

The removal and installation of storage devices, such as hard drives and optical drives, is pretty straightforward. There really isn't any deviation in the process of installing or exchanging the hardware. Fortunately, with today's operating systems, little to no configuration is required for such devices. The Plug and Play BIOS and operating system work together to recognize the devices. However, you still have to partition and format out-of-the-box hard drives before they will allow the installation of the operating system. Nevertheless, today's operating systems allow for a pain-free partition/format/setup experience by handling the entire process, if you let them.

Removing Storage Devices

Removing any component is frequently easier than installing the same part. Consider the fact that most people could destroy a house, perhaps not safely enough to ensure their well-being, but they don't have to know the intricacies of construction to start smashing away. On the other hand, very few people are capable of building a house. Similarly, many could

figure out how to remove a storage device, as long as they can get into the case to begin with, but only a few could start from scratch and successfully install one without tutelage.

This section details the removal of internal storage devices, and the section "Installing Storage Devices" details their installation. The focus here is on internal devices as opposed to external ones. Today's external storage devices are eSATA-, USB-, and perhaps Thunderbolt-attached devices, making them completely Plug and Play. Only the software preparation of external hard drives is a consideration, but the same procedure for the software preparation of internal devices works for external devices as well.

In Exercise 2.2, you'll remove an internal storage device.

EXERCISE 2.2

Removing an Internal Storage Device

1. With the power source removed from the system, ground yourself and the computer to the same source of ground.

   The most reliable way to ground yourself is to turn the computer off but leave it plugged into an electrical outlet. Then, wear an antistatic wrist strap and attach the alligator clip end to an unpainted metal portion of the case you're working in. This will effectively drain static electricity from your body before you touch the sensitive electronic components. If you don't have an antistatic strap, touch and keep touching (or frequently touch) an unpainted metal portion of the case to remove excess static electricity.

2. Remove the cover from the system, exposing the internal components.
3. Unplug all connections from the storage device you wish to remove. These include data and power connections as well as any others, such as audio connections to the sound card or motherboard (for an optical drive). Always be sure to grip the connector, not the wires.
4. Gather the appropriate antistatic packaging for all static-sensitive components that will be reused in the future, including any adapter cards that the storage device plugs into.
5. Remove any obstructions that might hinder device removal, such as component cables attached to adapter cards or the adapter cards themselves, storing them in antistatic packaging so that they can be reused.
6. Remove related adapter cards from the motherboard, storing them in antistatic packaging so that they can be reused.
7. Remove the machine screws holding the storage device to the chassis. These could be on the side of the device or on the bottom.

EXERCISE 2.2 *(continued)*

8. Some devices, especially hard drives because they have no front access from the case, pull out of the chassis toward the rear of the case, whereas others, such as optical drives, generally pull out from the front. A gentle nudge from the rear of the device starts it on its way out the front. Go ahead and remove the device from the case. If you discover other components that obstruct the storage device's removal, repeat step 5.

Installing Storage Devices

An obvious difference among storage devices is their *form factor*. This is the term used to describe the physical dimensions of a storage device. Form factors commonly have the following characteristics:

- 3½″ wide vs. 5¼″ wide
- Half height vs. full height vs. 1″ high and more
- Any of the laptop specialty form factors

You will need to determine whether you have an open bay in the chassis to accommodate the form factor of the storage device that you want to install. Adapters exist that allow a device of small size to fit into a larger bay. For obvious reasons, the converse is not also true.

In Exercise 2.3, you'll install an internal storage device.

EXERCISE 2.3

Installing an Internal Storage Device

1. With the power source removed from the system, ground yourself and the computer to the same source of ground.
2. Remove the cover from the system, exposing the internal components.
3. Locate an available bay for your component, paying attention to your device's need for front access. If you do not see one, look around; some cases provide fastening points near the power supply or other open areas of the case. If you still do not see one, investigate the possibility of sacrificing a rarely or never used device to make room.
4. Remove any obstructions that might hinder device installation, such as component cables attached to adapter cards or the adapter cards themselves, storing them in antistatic packaging to be reused.
5. Find the proper screws for the storage device, and set any jumpers on the drive while it is in hand. Then insert the device into the bay. Keep in mind that some insert from the rear of the bay and some from the front.
6. Line up the screw holes in the device with the holes in the bay. Note that many devices rarely insert as far as they can before lining up with the chassis's holes, so don't be

surprised when pushing the device all the way into the bay results in misalignment. Other devices that require front access stop themselves flush with the front of the case, and still others require you to secure them while holding them flush.

7. Use at least two screws on one side of the device. This keeps the device from sliding in the bay as well as from rotating, which happens when you use only one screw or one screw on each side. If the opposite side is accessible, go ahead and put at least one screw in the other side. Most devices allow for as many as four screws per side, but eight screws are not necessary in the vast majority of situations.
8. Connect the data cable from the device to the adapter card or motherboard header.
9. Attach a power connector from the power supply to the device, if needed (it is for PATA, SATA, and SAS). Be sure to insert the connector completely.
10. Once the drive is attached, unground yourself, and turn the computer on to verify that the drive is functional.
11. If the drive is working properly, replace the computer cover.

Understanding Power Supplies

The computer's components would not be able to operate without power. The device in the computer that provides this power is the *power supply* (see Figure 2.26). A power supply converts 100–240V AC current into the DC voltages that a computer needs to operate. These are +3.3VDC, +5VDC, +12VDC, and –12VDC. The jacket on the leads carrying each type of voltage has a different industry-standard color-coding for faster recognition.

FIGURE 2.26 A desktop power supply

You will see the term PSU in reference to a *power supply unit*.

Throughout this section, you will see us use the terms *watts*, *volts*, and *amps*. If you're working with electricity a lot, you might also see the term *ohms*. To help understand what these terms mean, let's use an analogy of water flowing through a pipe. Amps would be the amount of water flowing through the pipe; voltage would be the water pressure; and watts would be the power that the water could provide. (Watts mathematically are volts × amps.) If there were a filter or other barrier in the pipe, that would provide resistance, which is measured in ohms. In non-analogous terms, amps are the unit of current flow; volts are the unit of force; watts are the unit for power (watts = volts × amps); and ohms are resistance.

The abbreviation *VDC* stands for *volts DC. DC* is short for *direct current*. Unlike alternating current (AC), DC does not alter the direction in which the electrons flow. AC for standard power distribution does so 50 or 60 times per second (50 or 60 Hz, respectively). Sometimes you will see people abbreviate a current such as positive 5 volts as +5V, and other times as +5VDC. It's really a matter of preference.

Be aware that DC voltage is not safer than AC voltage, despite its common use in batteries and lower-power components. Direct current is more likely to cause a prolonged clamping of the muscles than AC, which is more likely to fibrillate the heart, which results in a deadly loss of coordination of the various cardiac muscles. Furthermore, power supplies contain transformers and capacitors that can discharge *lethal* amounts of current even when disconnected from the wall outlet for long periods. They are not meant to be serviced, especially by untrained personnel. Do *not* attempt to open them or do any work on them. Simply replace and recycle them when they go bad.

A+ Exam Objective 3.6, "Given a scenario, install the appropriate power supply," highlights input of 110-120VAC and 220-240VAC, as well as output of 3.3V vs. 5V vs. 12V. Understand the voltages power supplies require and produce, including wattage ratings.

Power Supply Input

Computer power supplies need to get their power from somewhere, and that is typically a wall outlet. There may be an intermediary battery backup device in-between called an uninterruptible power supply (UPS), which we will talk about later in the "Battery Backup Systems" section, but the point is that the power supply doesn't just generate its own power. It converts AC power from the wall into DC power that components use.

Countries have differing standards on the voltage provided by wall outlets. In the United States, it's typically 110V and 220V. The 110V outlets are the "normal" outlets that most electronics, including computers, are plugged into. The 220V outlets are for high-energy devices such as electric ranges and clothes dryers. Fortunately, the two plugs are completely different (as shown in Figure 2.27) to help us avoid plugging the wrong thing into the wrong place and frying the component. As noted, though, other countries have different standards, and power supply manufacturers want to ensure their devices work in different countries.

FIGURE 2.27 110V (left) and 220V (right) wall outlets

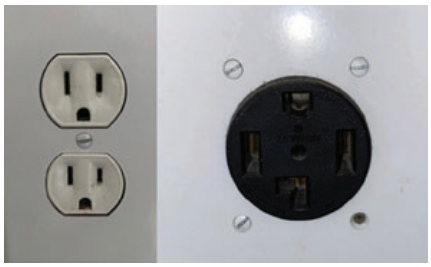

Therefore, some power supplies have a recessed, two-position slider switch, often a red one, on the rear that is exposed through the case. *Dual-voltage options* on such power supplies read 110 and 220, 115 and 230, or 120 and 240. This selector switch is used to adjust for the voltage level used in the country where the computer is in service. Newer PSUs will autodetect the incoming power level and adjust accordingly—they are referred to as *universal voltage* power supplies.

As noted earlier, in the United States, the power grid supplies anywhere from 110V to 120V. However, in Europe, for instance, the voltage supplied is double, ranging from 220V to 240V. Although the voltage is the same as what is used in the United States to power high-voltage appliances, the amperage is much lower. The point is, if the PSU has a switch, it's not there to allow multiple types of outlets to be used in the same country. If the wrong voltage is chosen in the United States, the power supply will expect more voltage than it receives and might not power up at all. If the wrong voltage is selected in Europe, however, the power supply will receive more voltage than it is set for. The result could be disastrous for the entire computer and could result in sparking or starting a fire. Always check the switch before powering up a new or recently relocated computer. In the United States and other countries that use the same voltage, check the setting of this switch if the computer fails to power up. Or, simply use universal voltage power supplies.

Power Supply Output and Wattage Ratings

Power supplies all provide the same voltages to a system, such as +3.3V, +5V, +12V, and –12V. Each of these can be referred to as a *rail*, because each one comes from a specific tap (or rail) within the power supply. Some power supplies provide multiple 12V rails in an effort to supply more power overall to components that require 12V. For instance,

in *dual-rail power supplies*, one rail might be dedicated to the CPU, whereas the other is used to supply power to all of the other components that need 12V.

The problem that can arise in high-powered systems is that although the collective power supplied by all rails is greater than that supplied by power supplies with a single rail, each rail provides less power on its own. As a result, it is easier to overdraw one of the multiple rails in such a system, causing a protective shutdown of the power supply. Care must be taken to balance the load on each of the rails if a total amperage greater than any one rail is to be supplied to attached components. Otherwise, if the total power required is less than any single rail can provide, there is no danger in overloading any one rail.

Power supplies are rated in watts. A watt is a unit of power. The higher the number, the more power your computer can draw from the power supply. Each power supply comes with a *wattage rating*—think of this rating as the "capacity" of the device to supply power. Most computers require power supplies in the 500- to 600-watt range. Higher wattage power supplies, say 800- to 1300-watt, might be required for more advanced systems that employ power-hungry graphics cards or multiple disk drives, for instance. As of this writing, power supplies of up to 2,000 watts were available for desktop machines. It is important to consider the draw that the various components and subcomponents of your computer place on the power supply before choosing one or its replacement.

Real World Scenario

How Much Power Do You Really Need?

Having a power supply that doesn't provide sufficient wattage to system components will result in brownouts, intermittent reboots, or even the system failing to power up. Having too much power is a little inefficient, but you won't have to worry about failures such as those. The question is, how much do you really need?

Most power-hungry devices (and we're looking at you, video cards) will give you a minimum power supply requirement. Generally speaking, the video card is the biggest power user, so if you meet these requirements, you should be fine. The manufacturer assumes you have otherwise standard gear in the system and adds a bit of a buffer.

Another option, if you are building a system, is to use the System Builder feature on a website we love called `PCPartPicker.com`. As you add components to the system, the site automatically calculates the estimated wattage drawn and displays it on the page. Pick a PSU that supplies more than that amount of power and you should be in the clear.

Energy Efficiency

As you've learned, power supplies take AC current from the wall and convert it into DC power for PC components. The conversion isn't 100 percent efficient, and heat generation is a side effect of the lack of efficiency. Of course, excess heat is an enemy of sensitive electronic

components. In addition, wasted energy conversion means that more energy is needed to produce the same output, which will increase electric bills.

Energy efficiency is part of Exam Objective 3.6. Understand how it works and what it means when selecting a power supply.

In 2004, the 80 PLUS certification program was launched to rate power supplies on their *energy efficiency*—that is, how much of the input energy they convert into output energy. PSUs that have 80 percent or greater efficiency at 20 percent, 50 percent, and 100 percent of their rated loads can qualify for 80 PLUS certification. There isn't a law or rule that requires manufacturers to certify their PSUs, nor are there requirements for users or companies to purchase certified PSUs. The certification is designed to help users make better choices for their computing situations, and potentially save money and wear-and-tear on their components in the process.

There are currently six standards of 80 PLUS certification. The higher the level, the more efficient the PSU is. Table 2.7 shows the levels and minimum efficiencies required to obtain the certification. It's highly unlikely that you will be tested on the efficiency thresholds, but it would be a good idea to know the relative rankings, just in case.

TABLE 2.7 80 PLUS rating minimum efficiency

Efficiency at	80 PLUS	80 PLUS Bronze	80 PLUS Silver	80 PLUS Gold	80 PLUS Platinum	80 PLUS Titanium
20% load	80%	82%	85%	87%	90%	90%
50% load	80%	85%	88%	90%	92%	92%
100% load	80%	82%	85%	87%	89%	94%

At a minimum, you should look for 80 PLUS–certified PSUs. Going higher than that depends on the circumstances. Users running powerful, high-end systems with energy-hungry components can save several dollars a month on energy bills by going with a 80 PLUS Gold or higher PSU. Over time, a more efficient PSU can save a user money, not to mention potentially extend the lifetime of the computer's components.

Power Connectors

The connectors coming from the power supply are quite varied these days. Some PSUs will have connectors permanently attached, whereas other PSUs give you the ability to attach and detach power connectors as needed, based on the devices installed in the system. The following sections detail and illustrate the most common power connectors in use today.

For Exam Objective 3.6, be familiar with the 20+4 pin motherboard power connector.

ATX and ATX12V Connectors

ATX motherboards use a single block connector from the power supply. When ATX boards were first introduced, this connector was enough to power all the motherboard, CPU, memory, and all expansion slots. The original ATX system connector provides the six voltages required, plus it delivers them all through one connector: an easy-to-use single 20-pin connector. Figure 2.28 shows an example of an ATX system connector.

FIGURE 2.28 20-pin ATX power connector

The development of PCIe was great for the computer industry but challenging from a power standpoint because PCIe devices require more power than PCI ones did. For ATX motherboards with PCIe slots, the 20-pin system connector proved inadequate. This led to the ATX12V 2.0 standard and the even higher-end EPS12V standard for servers. These specifications call for a 24-pin connector that adds further positive voltage leads directly to the system connector. The 24-pin connector looks like a larger version of the 20-pin connector. The corresponding pins of the 24-pin motherboard header are actually keyed to accept the 20-pin connector. Adapters are available if you find yourself with the wrong combination of motherboard and power supply. Some power supplies feature a 20-pin connector that snaps together with a separate 4-pin portion for flexibility, called a 20+4 connector, which can be seen in Figure 2.29. Otherwise, it will just have a 24-pin connector.

The adapter mentioned in the previous paragraph is called a *20-pin to 24-pin motherboard adapter*.

ATX12V 2.1 introduced a new 6-pin connector (Figure 2.30), which was specifically designed to give additional dedicated power to the PCIe adapters that required it. It provided a 75W power source to such devices. ATX12V 2.2 replaced the 75W 6-pin connector with a 150W 8-pin connector. The connector in Figure 2.30 is modular. It comes standard as a 6-pin connector, but also has a 2-pin additional attachment to make it into an 8-pin one.

FIGURE 2.29 A 24-pin ATX12V 2.x connector, in two parts

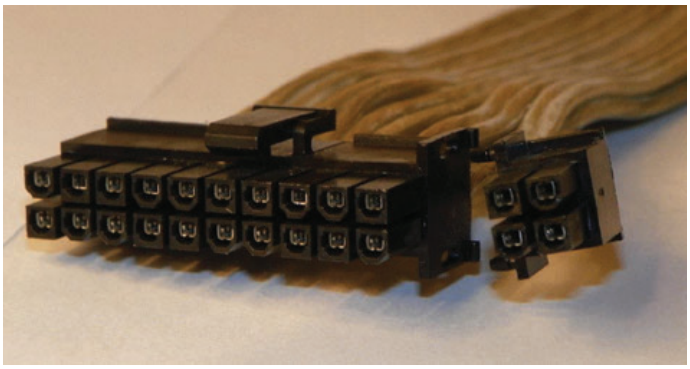

FIGURE 2.30 6-pin ATX12V 2.1 PCIe connector with extender for an 8-pin ATX12V 2.2 PCIe connection

SATA Power Connectors

SATA drives arrived on the market with their own power requirements in addition to their new data interfaces. (Refer back to Figure 2.12 to see the SATA data and power connectors.) For SATA drives, use the 15-pin SATA power connector, a variant of which is shown in Figure 2.31. The fully pinned connector is made up of three +3.3V, three +5V, and three +12V leads interleaved with two sets of three ground leads. Each of the five sets of three common pins is supplied by one of five single conductors coming from the power supply. When the optional 3.3V lead is supplied, it is common to see it delivered on an orange conductor.

Older PATA hard drives and optical drives used a thicker 4-pin power connector called a Molex connector. An example is shown in Figure 2.32. The connector is keyed with two beveled corners, making it nearly impossible to install incorrectly.

FIGURE 2.31 SATA power connector

FIGURE 2.32 Molex power connector

Modular Power Supplies

On older PSUs, all power connectors were hardwired into the power supply itself. This so-called *fixed cable* PSU had a number of interesting side effects. One was that no matter how many or how few internal devices were present, there were a fixed number of connectors. Power supply manufacturers generally provided enough so most users wouldn't run short of connectors. The flip side of this was that there were often four to six unused connectors, but the cables were still taking up space inside the case. Zip ties and thick rubber bands helped maintain the chaos.

Modular power supplies are part of Exam Objective 3.6. Know the difference between them and fixed cable PSUs, and how to install them.

As the variety of internal components became more complex, the need arose to have more flexibility in terms of the connectors provided. Out of this need rose an elegant solution—the *modular power supply*. From a functional standpoint, it works just as a non-modular power

supply does. The difference is that none of the power cables are permanently attached—only the ones that are needed are connected. The PSU shown back in Figure 2.26 shows the side of a fully modular power supply. The top row has connectors for the motherboard (left and center) and CPU or PCIe device. On the bottom row, you can see four 6-pin peripheral connectors and three 8-pin ones to power the CPU or PCIe devices. Nearly all PSUs sold in the market today are modular.

You will also occasionally see semi-modular PSUs on the market. Generally, the motherboard and CPU connectors will be hardwired, whereas the peripheral connectors can be added as needed. There are two potential disadvantages to using a fully modular or semi-modular power supply. First, some PSU manufacturers use proprietary connectors. Always be sure to keep the extra power connectors around (many come with a bag to store unused cables) just in case they are needed. Second, modular PSUs can take up a little more room in the case. Plugging the power connectors into the PSU can take up to an extra half inch. Usually this isn't an issue, but it can be in smaller cases.

Redundant Power Supplies

Nearly every computer you will work with has one and only one power supply—is that enough? If the PSU supplies the right amount of wattage to safely power all components, then the answer is nearly always yes. There are some instances, though, where power redundancy is helpful or even critical. Within the realm of power redundancy, there are two paths you can take: redundant power supplies within a system or battery backups. Let's look at both.

Redundant power supplies are included in Exam Objective 3.6. Understand how they work and how to install one.

Multiple PSUs

It's almost unheard of to see two power supplies installed in a desktop computer. There's generally no need for such a setup, and it would just be a waste of money. And for laptops and mobile devices, it's simply not an option. For servers, though, having a *redundant power supply (RPS)*, meaning a second PSU installed in the system, might make sense. The sole reason to have two power supplies is in case one fails, the other can take over. The transition between the two is designed to be seamless and service will not be disrupted.

Based on its name and our description so far, it might seem as though this means installing two full-sized PSUs into a computer case. Given the limited amount inside a case, you can imagine how problematic this could be. Fortunately, though, PSU manufacturers make devices that have two identical PSUs in one enclosure. One such example is shown in Figure 2.33. The total device is designed to fit into ATX cases and is compliant with ATX12V and EPS12V standards. If one fails, the other automatically takes over. They are hot-swappable, so the failed unit can be replaced without powering the system down.

FIGURE 2.33 Hot-swappable redundant PSUs

Source: Rainer Knäpper / Wimimedia Commons / CC BY-SA 2.0 DE

Although an RPS can help in the event of a PSU failure, it can't keep the system up and running if there is a power outage.

Battery Backup Systems

The second type of power redundancy is a battery backup system that the computer plugs into. This is commonly referred to as an *uninterruptible power supply (UPS)*.

These devices range in size from as small as a brick to as large as a room. Some just have a few indicator lights, whereas others have LCD displays that show status and menus and come with their own management software. The UPS will have several power plugs. It might divide the plugs such that a few of them provide surge protection only, whereas others provide surge protection as well as backup power. Figure 2.34 shows a small UPS that could be used in a home or small office to protect a few machines, and Figure 2.35 shows a room-sized one that can support dozens of servers in the event of a power failure. It's approximately 10 feet long, 6 feet wide, and 6 feet tall.

FIGURE 2.34 An uninterruptible power supply front (a) and back (b)

(a) (b)

FIGURE 2.35 Large UPS

Inside the UPS are one or more batteries and fuses. Much like a surge suppressor, a UPS is designed to protect everything that's plugged into it from power surges. UPSs are also

designed to protect against power sags and even power outages. Energy is stored in the batteries, and if the power fails, the batteries can power the computer for a period of time so that the administrator can then safely power it down. Many UPSs and operating systems will also work together to safely power down automatically a system that gets switched to UPS power. These types of devices may be overkill for Uncle Bob's machine at home, but they're critically important fixtures in server rooms.

The UPS should be checked periodically to make sure that its battery is operational. Most UPSs have a test button that you can press to simulate a power outage. You will find that batteries wear out over time, and you should replace the battery in the UPS every couple of years to keep the UPS dependable.

UPSs all have a limit as to how many devices they can handle at once. There are three numbers to pay attention to when looking for a UPS:

- Volt-ampere (VA) rating
- Power factor (PF)
- Watts

The VA rating shows the maximum load the UPS can support. The higher the number, the bigger the load. A UPS will also have a power factor, which is a measure of efficiency, somewhere between 0 and 1.0. Again, the higher, the better. Finally, the wattage supported is calculated by taking the VA × PF. For example, a 1,000 VA load UPS with a 0.8 PF will support a maximum of 800 watts. Always choose a UPS with a VA of at least 1.2 times the anticipated need. UPS power limitations should be strictly observed. If overloaded, it can cause a short, which could potentially result in fire.

Replacing Power Supplies

Sometimes power supplies fail. Sometimes you grow out of your power supply and require more wattage than it can provide. Often, it is just as cost effective to buy a whole new case with the power supply included rather than dealing with the power supply alone. However, when you consider the fact that you must move everything from the old case to the new one, replacing the power supply becomes an attractive proposition. Doing so is not a difficult task.

Regardless of which path you choose, you must make sure the power connection of the power supply matches that of the motherboard to be used. Additionally, the physical size of the power supply should factor into your purchasing decision. If you buy a standard ATX-compatible power supply, it might not fit in the petite case you matched up to your micro ATX motherboard. In that scenario, you should be on the lookout for a smaller form factor power supply to fit the smaller case. Odds are that the offerings you find will tend to be a little lighter in the wattage department as well.

Exercise 2.4 details the process to remove an existing power supply. Use the reverse of this process to install the new power supply. Just keep in mind that you might need to procure the appropriate adapter if a power supply that matches your motherboard can no longer be found. There is no post-installation configuration for the power supply, so there is nothing to cover along those lines. Many power supply manufacturers have utilities on their websites that allow you to perform a presale configuration so that you are assured of obtaining the most appropriate power supply for your power requirements.

EXERCISE 2.4

Removing a Power Supply

1. With the power source removed from the system, ground yourself and the computer to the same source of ground.
2. Remove the cover from the system, exposing the internal components.
3. After locating the power supply, which can come in a variety of formats and appear on the left or right side of the case, follow all wiring harnesses from the power supply to their termini, disconnecting each one.
4. Remove any obstructions that appear as if they might hinder the removal of the power supply.
5. Using the dimensions of the power supply, detectable from the inside of the case, note which machine screws on the outside of the case correspond to the power supply. There are often four such screws in a non-square pattern. If your case has two side panels, and you removed only one, there will likely be one or more screws holding the other panel on that appear to be for the power supply. These do not need to be removed. If all case screws have been removed, pay attention to their location and do not use these holes when securing the new power supply.
6. Remove the screws that you identified as those that hold the power supply in place. Be aware that the power supply is not lightweight, so you should support it as you remove the final couple of screws.
7. Maneuver the power supply past any obstructions that did not have to be removed, and pull the power supply out of the case.

AC Adapters as Power Supplies

Just as the power supply in a desktop computer converts AC voltages to DC for the internal components to run on, the AC adapter of a laptop computer converts AC voltages to DC for the laptop's internal components. And AC adapters are rated in watts and selected for use with a specific voltage just as power supplies are rated. One difference is that AC adapters

are also rated in terms of DC volts out to the laptop or other device, such as certain brands and models of printer.

Because both power supplies and AC adapters go bad on occasion, you should replace them both and not attempt to repair them yourself. When replacing an AC adapter, be sure to match or exceed the wattage requirements for the device. In addition, while new AC adapters use USB-C, older AC adapters may have proprietary connectors and either positive or negative polarity. You can read more on this subject in Volume 1 Chapter 4, "Laptop and Mobile Device Hardware."

Summary

In this chapter, you learned about three classes of personal computer components that finish our tour of the inside of the case—expansion cards, storage devices, and power supplies.

Expansion cards add helpful capabilities, such as video, audio, network connections, and additional ports for devices and peripherals. Storage devices provide long-term data capacity. Examples include conventional spinning hard drives and SSDs. RAID arrays can help provide performance increases and additional data protection. Other removable storage devices include flash drives, memory cards, and optical drives.

Storage communications' interfaces you learned about included NVMe, SATA, PCIe, and SAS. Form factors covered included M.2 and mSATA.

Finally, we discussed power supply safety as well as the various requirements and connectors. You learned about how power supplies take AC current and convert it into DC power for internal components in 2.3V, 5V, and 12V. You also learned about redundant power supplies, modular power supplies, wattage ratings, and efficiency ratings. Finally, you learned how to remove, install, and configure storage devices and how to replace power supplies.

Exam Essentials

Know how to install and configure expansion cards to provide needed functionality. Understand the functionality that video cards, sound cards, network cards, and capture cards provide. Know where to install them and broadly how to configure them.

Be familiar with the components of a conventional hard drive and the anatomy of a hard drive. Conventional hard drives store data magnetically on rotating metal platters. Be familiar with magnetic hard drive speeds, including 5,400, 7,200, 10,000, and 15,000 rpm. Form factors are 2.5″ and 3.5″.

Understand the advantages that solid-state drives have over conventional drives. SSDs are much faster than magnetic hard drives, produce less heat, and can be made much smaller physically. They are also less susceptible to shock from drops.

Know the differences between three SSD communications interfaces and two form factors. The SSD communications interfaces are NVMe, SATA, PCIe, and SAS. The two form factors to know are M.2 and mSATA.

Understand the details surrounding optical storage. From capacities to speeds, you should know what the varieties of optical storage offer. Blu-ray is the newest, fastest, and largest capacity, followed by DVDs and finally CDs.

Understand the different flash drive and memory card options available. Know the differences between SD cards, CompactFlash, microSD, and miniSD. Be able to identify which cards can fit into specific types of slots natively or with an adapter.

Understand the characteristics of five types of RAID configurations. You need to know RAID 0, RAID 1, RAID 5, RAID 6, and RAID 10. RAID 0 is disk striping, which can improve speed but does not provide fault tolerance. RAID 1 is disk mirroring, which gives fault tolerance but no performance increase. RAID 5 is striping with parity, which can give some performance boost along with fault tolerance. RAID 6 is double-parity RAID. It's similar to RAID 5, with a second parity stripe. RAID 10, also known as RAID 1+0, adds mirroring to a striped set. Understand what hot-swappable means.

Know about power supplies and their connectors. Power supplies are commonly made in ATX and other, smaller form factors. Regardless of their type, power supplies must offer connectors for motherboards and internal devices. Know the differences among the connectors and why you might need a 20-pin to 24-pin motherboard adapter. Also understand why AC adapters are related to power supplies.

Understand power supply characteristics that determine performance. Power supplies can take AC input from 110V to 240V and may have a switch on the back to determine which source to expect. Output to internal components will be 3.3V, 5V, and 12V. Capacity is measured in watts and expressed as a wattage rating. The amount of AC energy converted to DC power is called energy efficiency—the higher, the better. There are also redundant and modular power supplies.

Know how to remove, install, and configure storage devices. Know the difference between the data and power connectors used on storage devices. Know what it means to partition and format a hard drive. Be aware of the physical differences in storage device form factors.

Know how to remove, install, and configure power supplies. Know the difference between the modern motherboard power headers and know when an adapter might be required. Be familiar with how to fasten power supplies to the chassis as well as how to unfasten them.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. You are installing a new graphics adapter in a Windows 11 system. Which of the following expansion slots is designed for high-speed, 3D graphics adapters?
   - **A.** USB
   - **B.** NVMe
   - **C.** PCI
   - **D.** PCIe

2. You need to install a new M.2 hard drive into a desktop computer. The motherboard has an open M.2 slot that is M-keyed. Which types of M.2 hard drives will you be able to install into that motherboard?
   - **A.** B-key
   - **B.** M-key
   - **C.** B+M-key
   - **D.** A, B, and C
   - **E.** B and C

3. You install a new NIC for a user, and he asks what it does. Which of the following best characterizes the functions of this device?
   - **A.** Multimedia
   - **B.** Communications
   - **C.** Input/Output
   - **D.** Storage

4. A client wants to ensure that their server does not lose power if their power supply fails. Which of the following best describes the type of power supply needed in this situation?
   - **A.** A power supply that provides twice as much wattage than what is needed to power system components
   - **B.** A power supply that provides a voltage output selector
   - **C.** A modular power supply
   - **D.** A redundant power supply

5. A client is looking for a desktop storage technology that provides the fastest possible data transfers. Which technology should you recommend?
   - **A.** SATA
   - **B.** NVMe
   - **C.** 10,000 rpm
   - **D.** 15,000 rpm

6. You are choosing an optical drive format for a client. The client needs to store video files which are about 20 GB in size. Which technology should they use?
   - **A.** Double-sided, double-layer DVD+R
   - **B.** Single-sided, single-layer Blu-ray disc
   - **C.** Double-sided, single-layer DVD-R
   - **D.** Double-sided, single-layer DVD+R

7. Which of the following best describes the concept of hot-swappable devices?
   - **A.** Power does not need to be turned off before the device is inserted or removed.
   - **B.** The device can be removed with power applied after it is properly stopped in the operating system.
   - **C.** Care must be taken when swapping the device because it can be hot to the touch.
   - **D.** The device can be swapped while still hot, immediately after powering down the system.

8. Of the following voltage pairings, which one accurately represents the input and output, respectively, of power supplies and AC adapters?
   - **A.** AC in, AC out
   - **B.** DC in, DC out
   - **C.** AC in, DC out
   - **D.** DC in, AC out

9. What are the output voltages that are commonly produced by PC power supplies? (Choose all that apply.)
   - **A.** +3.3VDC
   - **B.** +4.3VDC
   - **C.** +5VDC
   - **D.** +8VDC
   - **E.** +12VDC
   - **F.** –12VDC
   - **G.** +110VDC
   - **H.** –110VDC

10. You are installing a new power supply in a desktop computer. Which of the following statements about power supplies is true?
    - **A.** You must make sure that the voltage selector switch on the back of the power supply is switched to the lower setting if the computer is going to be used in Europe.
    - **B.** SATA hard drives most often use the same type of power connector that PATA hard drives use.
    - **C.** Power supplies require a 20-pin to 24-pin motherboard adapter to supply power to ATX-based motherboards.
    - **D.** Power supplies convert AC input to DC output.

11. You have been asked to install a new hard drive for a client. Which of the following is not a consideration when installing an internal storage device?

- **A.** You should match the form factor of the drive or adapt it to an available drive bay or slot.
- **B.** You should secure the drive with at least two screws on one side and preferably two on each side.
- **C.** Due to the high revolutions at which modern hard drives spin, you must secure an external power source because the internal power supplies do not have the capacity.
- **D.** You need to be sure that the routing of the drive's data cable, if applicable, does not obstruct the engineered flow of air across internal components.

12. A client just bought a new digital camera for his photography business. Which of the following is specifically designed as an internal storage device for cameras?

- **A.** Memory cards
- **B.** Optical discs
- **C.** Flash drives
- **D.** mSATA

13. You are installing a second hard drive into a desktop computer. The hard drive is an NVMe PCIe card. Which of the following steps will you need to take during installation?

- **A.** Install the hard drive into an open PCIe slot.
- **B.** Install the hard drive into an open PCIe slot. Connect the NVMe PCIe data cable.
- **C.** Install the hard drive into an open PCIe slot. Connect the NVMe PCIe power connector.
- **D.** Install the hard drive into an open PCIe slot. Connect the NVMe PCIe data cable and power connector.
- **E.** Install the hard drive into an open NVMe slot.
- **F.** Install the hard drive into an open NVMe slot. Connect the NVMe PCIe data cable.
- **G.** Install the hard drive into an open NVMe slot. Connect the NVMe PCIe power connector.
- **H.** Install the hard drive into an open NVMe slot. Connect the NVMe PCIe data cable and power connector.

14. You are purchasing a new video card in a desktop computer. For the best performance, which type of video card should you purchase?

- **A.** PCI x16
- **B.** PCI x128
- **C.** AGP
- **D.** PCIe x128
- **E.** PCIe x16

15. Which of the following is not a consideration when upgrading power supplies?
    - **A.** You might find that you do not have a matching motherboard connector on your new power supply.
    - **B.** You might find that your case has a nonremovable power supply.
    - **C.** You might find that your power rating is not adequate on the new power supply.
    - **D.** You might find that you do not have enough of the appropriate connectors coming from the power supply for the devices that you have installed.

16. You are choosing a flash standard for a client. They want the smallest flash cards possible. Which standard should you recommend?
    - **A.** SD
    - **B.** CompactFlash
    - **C.** microSD
    - **D.** miniSD

17. Your goal is to build the fastest computer possible. Which of the following hard drive technologies should you choose for the fastest read performance?
    - **A.** M.2 PCIe
    - **B.** M.2 SATA
    - **C.** NVMe M.2
    - **D.** NVMe SATA

18. You have been asked to configure a RAID 5 system for a client. Which of the following statements about RAID 5 is true?
    - **A.** It provides fault tolerance but no performance improvement.
    - **B.** It provides performance improvement but no fault tolerance.
    - **C.** If multiple drives in the array fail, the data can be saved provided that one drive is still operational.
    - **D.** It requires at least three drives to implement.

19. You have a motherboard with integrated video, and the video fails. You install a video add-on card. Which of the following statements are the most accurate? (Choose two.)
    - **A.** The add-on card will not function properly because of the integrated video.
    - **B.** The add-on card will function if it is bridged to the integrated video chipset.
    - **C.** The add-on card will be detected automatically and function properly if it's PCIe.
    - **D.** For the card to work properly, you might need to disable the integrated video in the BIOS.
    - **E.** The add-on card will work properly by using the integrated card's video driver.

20. When replacing a power supply, which of the following tends to vary among power supplies and must be chosen properly to support all connected devices?

A. Wattage

B. Voltage

C. Amperage

D. Resistance

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answer compares to the authors', refer to Appendix B.

Detail the process for removing a power supply from a computer chassis.

Chapter 3

Peripherals, Cables, and Connectors

THE FOLLOWING COMPTIA A+ 220-1201 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

✓ **3.0 Hardware**

✓ **3.1 Compare and contrast display components and attributes.**

- Types
  - Liquid crystal display (LCD)
    - In-plane switching (IPS)
    - Twisted nematic (TN)
    - Vertical alignment (VA)
  - Organic light-emitting diode (OLED)
  - Mini light-emitting diode (Mini-LED)
- Touch screen/digitizer
- Inverter
- Attributes
  - Pixel density
  - Refresh rates
  - Screen resolution
  - Color gamut

✓ **3.2 Summarize basic cable types and their connectors, features, and purposes.**

- Peripheral cables
  - USB 2.0
  - USB 3.0

    - Serial
    - Thunderbolt
- Video cables
    - High-definition Multimedia Interface (HDMI)
    - DisplayPort
    - Digital Visal Interface (DVI)
    - Video Graphics Array (VGA)
    - USB-C
- Hard drive cables
    - Serial Advanced Technology Attachment (SATA)
    - External SATA (eSATA)
- Adapters
- Connector types
    - MicroUSB
    - MiniUSB
    - USB-C
    - Molex
    - Lightning
    - DB9

Thus far, our discussion of computer components has focused primarily on those inside the case. With knowledge of the key internal components under your belt, it is time to turn our attention to the outside of the computer. Dozens of external devices are available to enhance a computer's functionality. We'll cover a variety of them that add video, audio, input, output, and storage capabilities.

Of course, to connect external peripherals, we need some sort of cable and connector. Not everything is wireless yet! Consequently, we will also discuss the interfaces and cables associated with common peripherals. With that, it's now time to think outside the box.

Understanding Peripheral Devices

Peripheral devices add much needed functionality to computers, beyond the core components. Having a fast processor and terabytes of hard drive space is great, but it doesn't complete the picture. Users need the ability to input data and easily see and use the output that the processor generates. Of course, the types of devices that can input or receive output are quite varied. In the following sections, we are going to break peripheral devices into four categories:

- Video
- Audio
- Input and output
- Storage

We realize that video and audio are indeed input or output devices, but because they are more specialized, we will cover them separately. After this section, you will have a good understanding of purposes of and uses for several common peripheral devices, as well as how they connect to a PC.

Knowledge of a variety of peripheral devices was previously tested on A+ exams. With the newest version, the objectives covering peripherals were removed. Even though you might not see these items on the exam, we feel it's important to cover them, as having this knowledge will make you a more complete technician. Besides, it will help you understand what all the cables and connectors are for!

Video Devices

The primary method of getting information out of a computer is to use a computer video display. Display systems convert computer signals into text and pictures and display them on a TV-like screen. As a matter of fact, early personal computers used television screens because it was simpler to use an existing display technology than to develop a new one. The most common video device used is a monitor.

Monitors

Most display systems work the same way. First, the computer sends a signal to a device called the *video card*—an expansion board installed in an expansion bus slot or the equivalent circuitry integrated into the motherboard—telling it to display a particular graphic or character. The video card then renders the character for the display—that is, it converts the single instruction into several instructions that tell the display device how to draw the graphic and sends the instructions to the display device based on the connection technology between the two. The primary differences after that are in the type of video card you are using (digital or analog) and the type of display (LCD, LED, IPS, and so forth).

Types of Monitors

PC monitors today are generally based on some form of *liquid crystal display (LCD)* technology. First used with portable computers and then adapted to desktop monitors, LCDs are based on the concept that when an electrical current is passed through a semi-crystalline liquid, the crystals align themselves with the current. When transistors are combined with these liquid crystals, patterns can be formed. Patterns can then be combined to represent numbers or letters. LCDs are relatively lightweight and don't consume much power.

A+ Exam Objective 3.1, "Compare and contrast display components and attributes," expects you to know the difference between display types LCD, IPS, TN, VA, OLED, and Mini-LED. The LCD inverter is also listed as an objective.

Liquid crystals produce no light, so LCD monitors need a lighting source to display an image. Traditional LCDs use a fluorescent bulb, typically about 8″ long and just slimmer than a pencil, called a *backlight* to produce light. Backlights require a high-voltage, high-frequency energy that's not provided directly by a power supply, and in LCD monitors the energy is provided by a separate component called the *inverter*. The inverter is a small circuit board installed behind the LCD panel that takes DC current and inverts it to AC for the backlight. If you are having problems with flickering screens or dimness, it's more likely that the inverter is the problem, not the backlight itself.

There are two things to keep in mind if you are going to replace an inverter. First, they store and convert energy, which means they have the potential to discharge that energy. To an inexperienced technician, they can be dangerous. Second, make sure the replacement inverter was made to work with the LCD backlight that you have. If they weren't made for each other, you might have problems with a dim screen or poor display quality.

Inverters can discharge energy, which can cause severe injury to you. Be careful when working with them!

There are three popular variants of LCD monitors in use today: *twisted nematic (TN)*, *in-plane switching (IPS)*, and *vertical alignment (VA)*. The three differ in the ways they align the liquid crystals in the display. Let's take a look at each:

Twisted Nematic TN is the oldest of the three technologies. It has somewhat restricted viewing angles and not the best color reproduction. Colors may appear washed out or too blended together. On the bright side, they have very little lag and can handle high refresh rates such as 240 Hz with ease. (We will talk about refresh rates in the upcoming "Adjusting Display Settings" section.) That makes them a popular choice for competitive gamers. They're not the best choice for high-end video needs, but they are inexpensive, which also makes them good for everyday office use.

In-Plane Switching IPS followed TN into the market, and they have the best viewing angles and color reproduction of all LCD monitors. They usually have a bit more lag than TN monitors do, but the difference is minimal in higher-end models. IPS monitors are ideal for those seeking the best color experience, such as graphic designers and video artists. They are also the best choice for someone who wants to mount them vertically, like a lot of programmers do.

Vertical Alignment VA has the best contrast ratios of the three, meaning the difference between the dark and bright colors is the best. Color reproduction tends to be good but not quite as good as IPS, and VA monitors do have a bit more lag (but we're talking milliseconds here). It's more a jack-of-all-trades technology than the other two options.

Most LCDs today use a panel of *light-emitting diodes (LEDs)* instead, which consume less energy, run cooler, and live longer than fluorescent bulbs. Therefore, when you see a monitor advertised as an LED monitor, it's really an LCD monitor with LED backlighting.

Recent advancements in LED technology have enabled manufacturers to make smaller LED backlights and pack many more together. The resulting technology is called *Mini light-emitting diode (Mini-LED)*, and it produces better contrast ratios (darker darks and brighter brights) than a standard LED. This can be helpful in bright environments. They will also be more expensive than LED displays. Compared to OLEDs (covered next), they will have better brightness (but not as dark of darks), similar resolution, slightly inferior viewing angles and color saturation, but also be a little less expensive.

The final type of LED monitor we'll cover is an *organic light-emitting diode (OLED)* display. An organic light-emitting compound forms the heart of the OLED, and it is placed between a positive anode and a negative cathode, which produce a current that runs through the electroluminescent compound, causing it to emit light. An OLED, then, is the combination of the compound and the electrodes on each side of it. The electrode in the back of the OLED cell is usually opaque, allowing a rich black display when the OLED cell is not lit. The front electrode should be transparent, to allow the emission of light from the OLED.

If thin-film electrodes and a flexible compound are used to produce the OLEDs, an OLED display can be made flexible, which is not only cool, but it allows it to function in places where other display technologies could never work.

Because OLEDs create the image in an OLED display and supply the light source, there is no need for a backlight, so power consumption is less than it is in LCD or LED panels. Additionally, the contrast ratio of OLED displays exceeds that of LCD/LED panels, meaning that in darker surroundings, OLED displays produce better images than LCD/LED panels produce. The viewing angles are also the best. The only performance downside of OLEDs is that they can be susceptible to *burn-in*, which is a permanent discoloration or artifact on a monitor screen that happens after an image has been displayed for a long time. That said, OLED monitors are generally some of the highest-quality monitors you will find on the market today, and usually more expensive than LEDs or Mini-LEDs. OLED is commonly found in smaller devices such as smartphones as well.

Nearly all monitors in use today are flat-screen monitors and measure no more than a few inches deep. Occasionally, though, you will run into a monitor that's not flat and can be over a foot deep. These relics are based on old television technology called the cathode ray tube (CRT). CRT monitors are bulky and heavy, and almost certainly use an analog signal (and VGA connector) instead of digital.

Using Displays for Input

If you have ever used a smartphone, and it's likely you have, you're already familiar with the concept of using a display for input. The screen of a smartphone acts as both a display and input device, and such a component is referred to as a *touchscreen*. Of course, touchscreens aren't just for small mobile devices. Laptops commonly include them, and many desktop monitors have touchscreen capabilities as well.

Touchpads and digitizers are part of Exam Objective 3.1. Understand how they work!

There are two types of touchscreens. *Resistive touchscreens* use pressure from a stylus or finger to register the contact. *Capacitive touchscreens* use electrical conductivity, which makes the human finger the perfect stylus. Resistive touchscreens allow for more granularity because they can account for more sensors per square inch. Capacitive touchscreens are more responsive and better suited to multitouch applications (which can register multiple touches at the same time), but they generally don't work if the user is wearing gloves, which can block the finger's electrical current. If you need help remembering the difference, think of resistive screens as providing resistance or needing pressure. And for capacitive screens, think of a capacitor, which stores electricity (and capacitive touchscreens utilize conductivity).

Touchscreens are said to digitize (or make digital) physical input. It follows, then, that touchscreens are a form of *digitizer*, or a device that can be written or drawn on, and the content will be converted from analog input to digital images on the computer. When built into the display, they might be the glass of the display itself, or they might be implemented

as an overlay for the display. Other types of digitizers are not built into the screen and may be USB devices that look like notepads, or pads on the palm rest or next to the touchpad of a laptop.

Adjusting Display Settings

Although most monitors are automatically detected by the operating system and configured for the best quality that they and the graphics adapter support, sometimes manually changing display settings, such as for a new monitor or when adding a new adapter, becomes necessary. Let's start by defining a few important terms:

- Refresh rate
- Screen resolution
- Multiple displays

Each of these terms relates to settings available through the operating system by way of display option settings.

Refresh Rate The *refresh rate*—technically, the vertical scan frequency—specifies how many times in one second the image on the screen can be completely redrawn, if necessary. Refresh rate is specified in hertz (Hz), which indicates how many times per second the screen is refreshed.

For LCD screens, the refresh rate may or may not be adjustable. The lowest standard refresh rate is 60 Hz, but higher-end monitors will be in the 360 Hz to 480 Hz range.

Higher refresh rates translate to more fluid video motion. Think of the refresh rate as how often a check is made to see if each pixel has been altered by the source. If a pixel should change before the next refresh, the monitor is unable to display the change in that pixel. Therefore, for gaming and home-theater systems, higher refresh rates are an advantage.

While the refresh rate is selected for the monitor, the refresh rate you select must be supported by both your graphics adapter and your monitor because the adapter drives the monitor. If a monitor supports only one refresh rate, it does not matter how many different rates your adapter supports—without overriding the defaults, you will be able to choose only the one common refresh rate. It is important to note that on some monitor and video card combinations, as the resolution you select increases, the higher supported refresh rates begin to disappear from the selection menu. If you want a higher refresh rate, you might have to compromise by choosing a lower resolution. Exercise 3.1 shows you where to change the refresh rate in Windows 11.

EXERCISE 3.1

Changing the Refresh Rate in Windows 11

1. Right-click a blank portion of the desktop.
2. Click Display Settings to open the System ➢ Display app, as shown in Figure 3.1.

EXERCISE 3.1 *(continued)*

FIGURE 3.1 Display settings

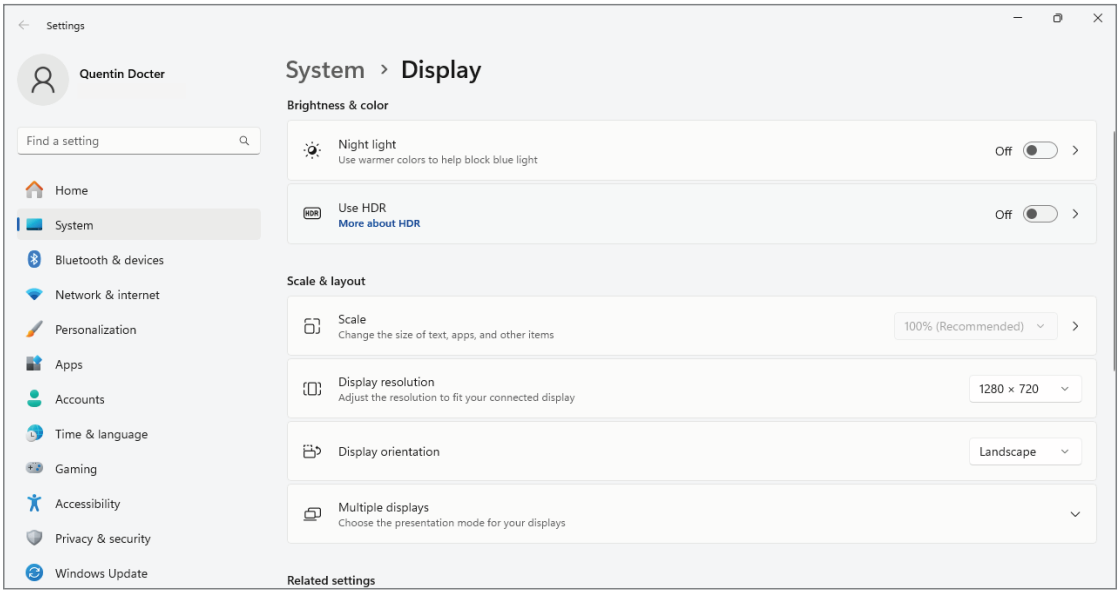

3. Near the bottom of the window in the Related Settings section, click Advanced Display. You will get the Advanced Display Settings window. For the computer in this example, only one monitor is attached. If multiple displays are present, the refresh rate can be independently set. The display choice is made using the drop-down menu at the top of the Advanced Display app (you can see it in Figure 3.2).

FIGURE 3.2 Advanced display refresh rates

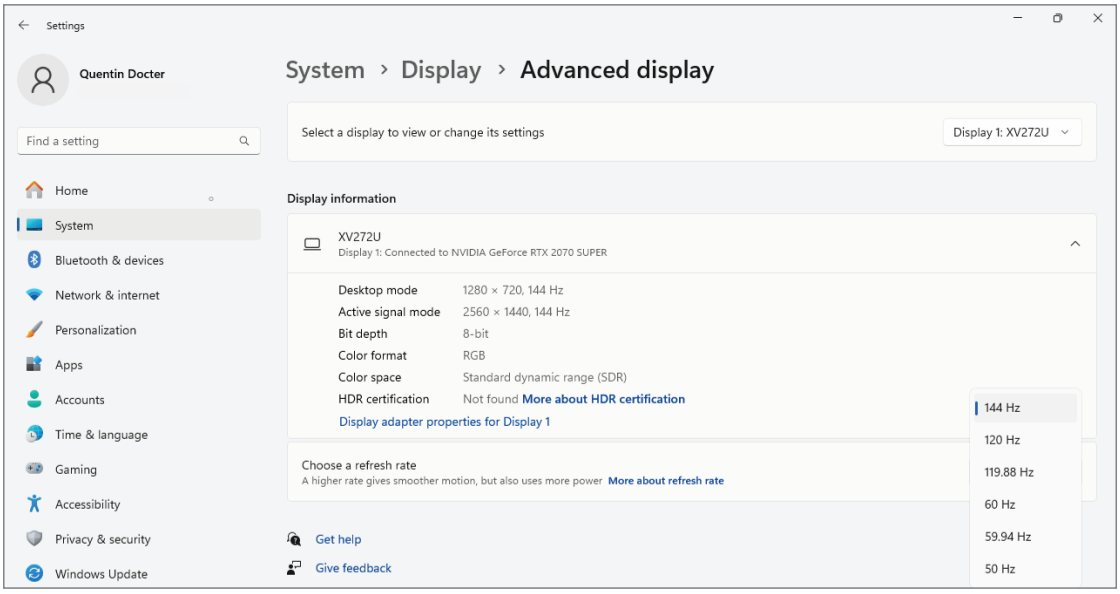

4. In the Choose A Refresh Rate box, the current refresh rate will be shown on the right side. Click the down arrow next to the refresh rate, and you will be able to see other refresh rate options, like in Figure 3.2. Note that many monitors will only have one available refresh rate.

   For this monitor running at 1280 × 720, there are six options. Note that if the resolution is set to the default 2560 × 1440, there are only three: 144 Hz, 120 Hz, and 59.95 Hz.

5. Choose the refresh rate you want for this selected monitor.

6. The display will go dark for a few seconds, and then your screen will reappear. You will have a pop-up window asking if you want to keep the changes or revert to previous settings. The default is to revert.

   Chances are you won't notice a difference after selecting a new refresh rate. The good news is that if it does mess up the image on the screen, you can wait about 15 seconds and the computer will revert back to the previous selection.

7. Close the Advanced Display window.

Eyestrain and Headaches

Some users (including this author) may complain about getting headaches when using displays with refresh rates set at 60 Hz or lower. The display screen flickers whenever it's refreshed, and it's almost always imperceptible.

As an aside, if you ever watch a live news broadcast and they are interviewing someone with a display behind them, you will often see the monitor flashing or what appears to be a rolling horizontal bar on the image on their screen. You're seeing the monitor refresh. Television and broadcast cameras are typically set to 30 Hz or 60 Hz (movie cameras usually operate at 24 Hz). If this is different than the display it's filming, it will cause the image to flicker or appear to roll.

Lower-frequency refresh rates such as 60 Hz can cause eyestrain, leading to headaches. In instances like this, invest in a monitor and video card that support higher refresh rates.

Screen Resolution *Resolution* is defined by how many software picture elements (pixels) are used to draw the screen. An advantage of higher resolutions is that more information can be displayed in the same screen area. A disadvantage is that the same objects and text displayed at a higher resolution appear smaller and might be harder to see. Up to a point, the added crispness of higher resolutions displayed on high-quality monitors compensates for the negative aspects.

The resolution is described in terms of the visible image's dimensions, which indicate how many rows and columns of pixels are used to draw the screen. For example, a resolution of 2560 × 1440 means 2560 pixels across (columns) and 1440 pixels down (rows) were used to draw the pixel matrix. The video technology in this example would use 2560 × 1440 = 3,686,400 pixels to draw the screen. Resolution is a software setting that is common among CRTs, LCDs, and projection systems, as well as other display devices.

Setting the resolution for your monitor is fairly straightforward. If you are using an LCD, for best results you should use the monitor's *native resolution*, which comes from the placement of the transistors in the hardware display matrix of the monitor. For a native resolution of 1920 × 1080, for example, there are 2,073,600 transistors (LCDs) or cells (OLED) arranged in a grid of 1920 columns and 1080 rows. Trying to display a resolution other than 1920 × 1080 through the operating system tends to result in the monitor interpolating the resolution to fit the differing number of software pixels to the 2,073,600 transistors, often resulting in a distortion of the image on the screen.

Some systems will scale the image to avoid distortion, but others will try to fill the screen with the image, resulting in distortion. On occasion, you might find that increasing the resolution beyond the native resolution results in the need to scroll the desktop in order to view other portions of it. In such instances, you cannot see the entire desktop all at the same time. The monitor has the last word in how the signal it receives from the adapter is displayed. Adjusting your display settings to those that are recommended for your monitor can alleviate this scrolling effect.

To change the resolution in Windows 11, right-click the desktop and choose Display Settings (as in Exercise 3.1). There is a pull-down menu for Display Resolution (it can be seen back in Figure 3.1). Click it and choose the resolution you want. Just like when you changed the refresh rate, changing the resolution will bring up a pop-up window asking if you want to keep the changes or revert back to the previous setting. If the change messes up the image and makes it unreadable, wait about 15 seconds and it will revert back to the previous setting.

Understanding Aspect Ratios

The term *aspect ratio* refers to the relationship between the horizontal and vertical pixel counts that a monitor can display. For example, old square-ish CRTs were shaped to support a display that conformed to 4:3 ratios, such as 800 × 600 or 1024 × 768. If you divide the first number by 4 and multiply the result by 3, the product is equal to the second number. Additionally, if you divide the first number by the second number, the result is approximately 1.3, the same as 4 ÷ 3. Displays with a 16:10 aspect ratio have measurements that result in a dividend of 16 ÷10 = 1.6.

When LCD monitors first became popular, they had wider screens and most supported a 16:10 ratio. Because the ATSC (Advanced Television Systems Committee) standard for widescreen television aspect ratios is 16:9 (1.778), computer monitors are trending more toward this same aspect ratio. As a result, the once popular 1920 × 1200, 16:10 resolution is uncommon and the 1920 × 1080, 16:9 resolution is the widescreen default. If you have a monitor that supports one and you try to set it to the other, the image may look squished or stretched, or the monitor may not display at all.

Multiple Displays Whether regularly or just on occasion, you may find yourself in a position where you need to use two or more monitors on the same computer simultaneously. For example, you may need to work in multiple spreadsheets at the same time and having two monitors makes it much easier. Or, if you are giving a presentation and would like to have a presenter's view on your laptop's LCD but need to project a slideshow onto a screen, you might need to connect an external projector to the laptop. Simply connecting an external display device does not guarantee that it will be recognized and work automatically. You might need to change the settings to recognize the external device or adjust options such as the resolution or the device's virtual orientation with respect to the built-in display. Exercise 3.2 guides you through this process.

When you have dual displays, you have the option to extend your desktop onto a second monitor or to clone your desktop on the second monitor. To change the settings for multiple monitors in Windows 11, follow the steps in Exercise 3.2, after ensuring that you have a second monitor attached.

EXERCISE 3.2

Changing the Settings for Multiple Monitors

1. Right-click a blank portion of the desktop.
2. Click Display Settings to open the Display Settings window.

   If a second monitor has been detected, you will see a screen similar to the one shown in Figure 3.3. Otherwise, you will need to click the Detect button in the Multiple Displays section.

   Notice that the first monitor is highlighted. If you were to change settings such as scale, resolution, or orientation, it would affect the monitor that's highlighted.

EXERCISE 3.2 *(continued)*

FIGURE 3.3 Multiple displays detected

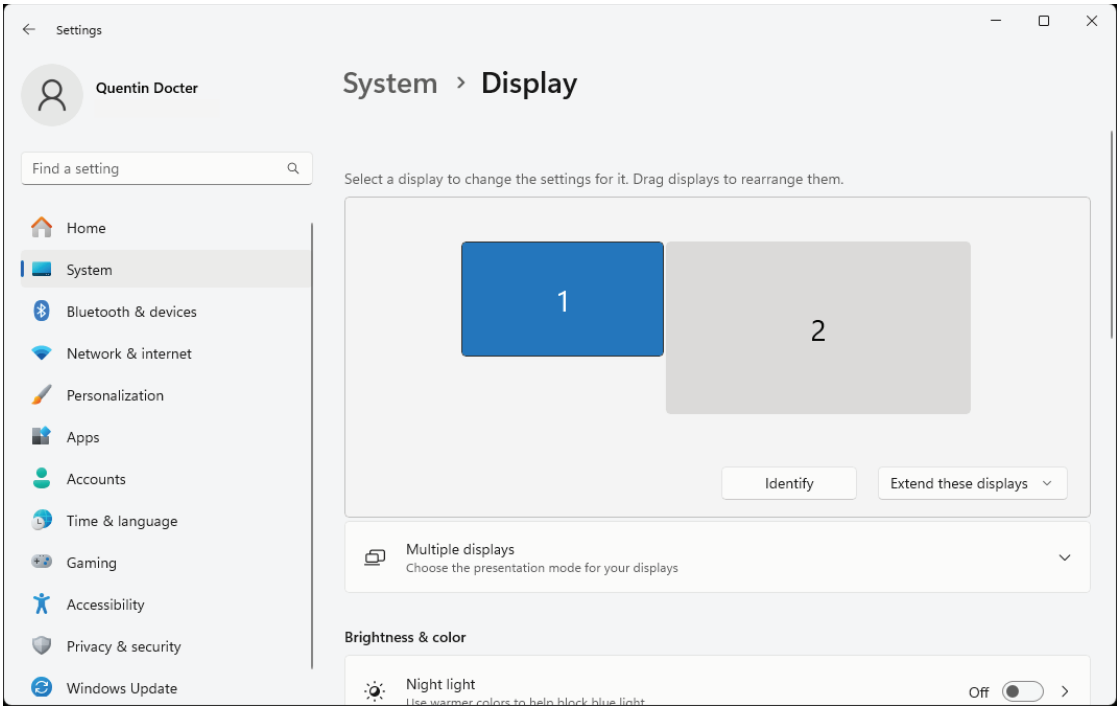

3. In the window showing the two displays, you have options to show an image on only one of the monitors, duplicate the displays, or extend the displays. In Figure 3.3, look where it says Extend These Displays, and click on the drop-down arrow. If your computer is not on Extend, choose that option now.
4. In the area showing two monitors, click and hold the second monitor and drag it around. Notice that you can place it above, below, left, or right of the first monitor. This will affect some display features, including where you need to move the mouse cursor to get it to appear on the second monitor.
5. Move the second monitor to be above the first monitor, and close Display Settings.
6. Move your mouse until you get the cursor to appear on the second screen.

   The rest of the steps are optional.
7. Open Display Settings and configure the second monitor to be in the position you want it to be relative to the first monitor.

   Does the second monitor (the one that's extended) show the Windows taskbar as well? If so, let's remove it. If not, let's add it.
8. Open the Taskbar settings app by right-clicking the taskbar and choosing Taskbar settings.

9. Scroll down to the Taskbar Behaviors section and click the drop-down arrow. In there, you will see the option Automatically Hide The Taskbar, as shown in Figure 3.4. If it's selected, deselect the box. If not, then select the box.

FIGURE 3.4 Taskbar settings

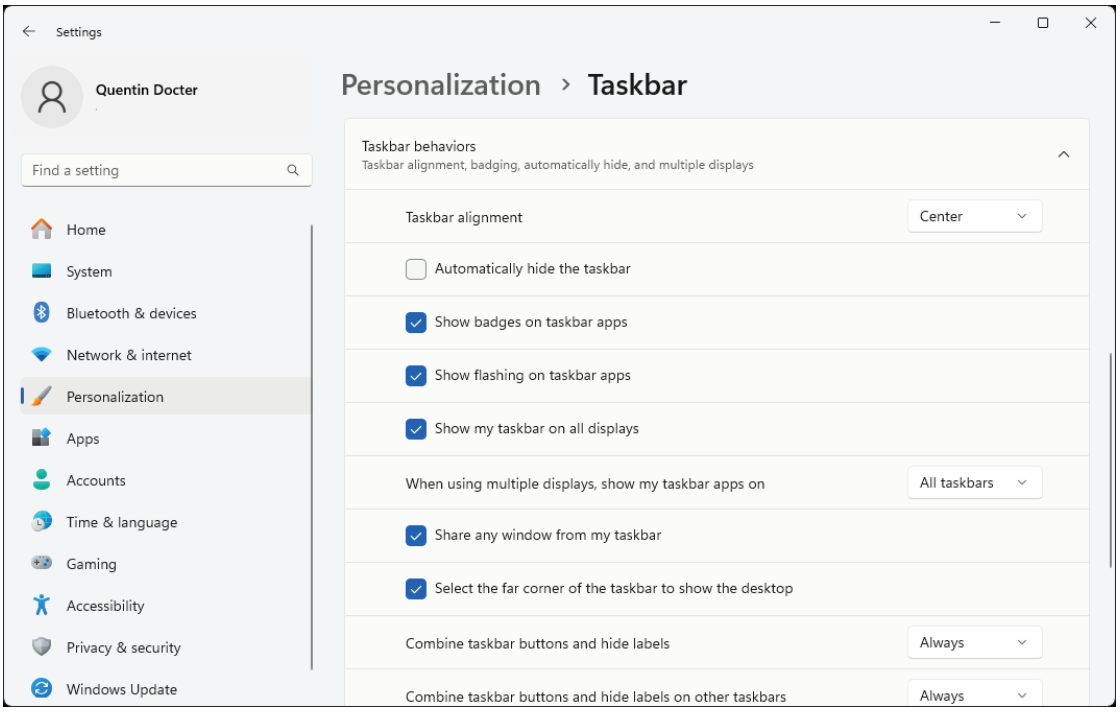

10. Configure the taskbar how you would like it, and then close the Taskbar app.

Selecting a Monitor

If you go to your favorite electronics retailer and search for monitors, the number of choices can be overwhelming. Here are a few tips to help narrow the field to a manageable number of options.

Exam Objective 3.1 lists four display attributes you need to know: pixel density, refresh rates, screen resolution, and color gamut.

Determine what the monitor will be used for. If it's a general-purpose monitor, it doesn't need to be the newest and fanciest technology. If it will be for a gamer, though, the minimum performance specifications just went up. For example, the resolution will need to be QHD or better and the refresh rate at least 144 Hz.

Choose a size. This depends on how far the user will be sitting from the screen and how much desk space is available. For a regular office or home office setup, somewhere

between 24″and 27″ is the most popular today, with some people preferring the slightly larger 32″. Anything bigger than that could be too large or not fit on the desk. If it will be viewed from across the room, then go as large as makes sense.

Pick a resolution. The resolutions available may be somewhat dependent on the size of monitor you choose. Remember that higher resolutions result in sharper images. Here are the three common options today:

1080p/Full High Definition (FHD) (1920 × 1080) This is the basic resolution used today. It's tenable for monitors up to 24″, but for larger monitors a higher resolution is needed; otherwise, icons and images appear oversized.

1440p/Quad High Definition (QHD) (2560 × 1440) For 27″ and larger displays, this is the base performance you want to look for. Sometimes you will hear this referred to as 2k resolution.

2160p/Ultra High Definition (UHD) (3840 × 2160) Commonly referred to as 4k, it's the resolution of choice for serious gamers and others who want ultra-sharp and crisp pictures. It's also great for those wanting to watch 4k movies on a monitor.

There are other resolutions available, but the three we listed are the most common. The highest available right now is 8k (UHD-2), which is 7680 × 4320. While you can buy 8k monitors, it may be too much of a good thing right now because of limited content available in 8k, and there are healthy debates on how much better it really is than 4k on screens smaller than 55″ due to limitations in human eyesight.

If you get an ultrawide monitor (and many are curved), you may see resolutions such as Wide QHD (WQHD [3440 × 1440]) or 4k offshoots like 7680 × 2160.

Choose a refresh rate. Basic, run-of-the-mill LCD screens will have a 60 Hz refresh rate; slightly nicer models will boast a 75 Hz rate. You probably won't notice too much of a difference between them. Gamers or those using their computer monitor to watch live-action sports or movies will want at least 144 Hz as a minimum. Hard-core gamers and ultimate performance seekers will want to go to at least 240 Hz. That kind of refresh rate is overkill for people staring at office applications and email all day.

Weigh optional features. Monitors today come with a wide variety of bells and whistles to enhance the user experience. Options to consider include curved or ultrawide vs. widescreen, number of I/O ports, integrated speakers and/or webcam and microphone, and voice assistant availability. Figure 3.5 shows the ports on a typical 27″ monitor. In the center are two HDMI ports, and moving to the right there is a DisplayPort, headphone jack, two USB-A ports, and a USB-B port. On the side of the monitor (not shown) are two additional USB-A ports.

FIGURE 3.5 Monitor ports

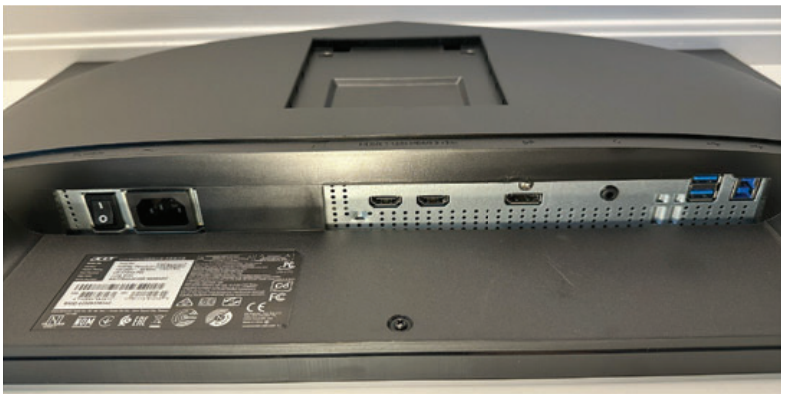

In addition, depending on the users' needs you might want to consider looking at *pixel density* or *color gamut* (both are test objectives, by the way). Pixel density refers to the number of pixels in an inch, or pixels per inch (PPI). The higher the number, the more pixels are in an inch, which can result in crisper and sharper images. Note that this is not the same as resolution. Resolution measures how many pixels wide and tall an image is. For example, imagine a QHD (2560 × 1440) image on a 24″monitor versus a 32″ monitor. Each one would use 2560 pixels wide, but the smaller monitor would need to have a greater PPI to squeeze them in.

Color gamut is the range of colors the display can produce. Most displays today can easily produce over 16 million colors, whereas the human eye can only detect around 1 million. So what's the big deal? A display with a higher color gamut can be important for visual designers and others who need high-quality color reproduction, because it will produce smoother and more vibrant images. Lower color gamut displays may produce washed-out looking images.

There are several color gamut standards. They include sRGB, which is standard for computer monitors and televisions; Adobe RGB, used for professional color imaging; NTSC, created for televisions; DCI-P3, used in movie theaters; and Rec. 709, used for high-definition TVs and monitors. It's unlikely that you'll need to memorize these standards any time soon, but do know that the higher the color gamut, the more colors the display can produce.

Consider the price. Last but not least is price. The higher the specs, the more it will cost. Balance the user's needs with their budget to determine the right monitor to select.

Projection Systems

Another major category of display device is the video projection system, or *projector*. A portable projector can be thought of as a condensed video display with a lighting system that projects the image onto a screen or other flat surface for group viewing. Interactive whiteboards have become popular over the past decade to allow presenters to project an image

onto the board as they use virtual markers to draw electronically on the displayed image. Remote participants can see the slide on their system as well as the markups made by the presenter. The presenter can see the same markups because the board transmits them to the computer to which the projector is attached, causing them to be displayed by the projector in real time.

To accommodate using portable units at variable distances from the projection surface, a focusing mechanism is included on the lens. Other adjustments, such as keystone, trapezoid, and pincushion, are provided through a menu system on many models as well as a way to rotate the image 180 degrees for ceiling-mount applications.

The key characteristics of projectors are resolution and brightness. Resolutions are similar to those of computer monitors. Brightness is measured in lumens. A *lumen (lm)* is a unit of measure for the total amount of visible light that the projector gives off, based solely on what the human eye can perceive and not on invisible wavelengths. Sometimes the brightness is even more of a selling point than the maximum resolution that the system supports because of the chosen environment in which it operates. For example, it takes a lot more to display a visible image in a well-lit office than it does in a darkened theater.

If you are able to completely control the lighting in the room where the projection system is used, producing little to no ambient light, a projector producing as little as 1,300 lumens is adequate in a home theater environment, while you would need one producing around 2,500 lumens in the office. However, if you can only get rid of most of the ambient light, such as by closing blinds and dimming overhead lights, the system should be able to produce 1,500 to 3,500 lumens in the home theater and 3,000 to 4,500 lumens in the office. If you have no control over a very well-lit area, you'll need 4,000 to 4,500 lumens in the home theater and 5,000 to 6,000 lumens in the business setting.

By way of comparison, a 60W standard light bulb produces about 800 lumens. Output is not linear, however, because a 100W light bulb produces over double, at 1,700 lm. Nevertheless, you couldn't get away with using a standard 100W incandescent bulb in a projector. The color production is not pure enough and constantly changes throughout its operation due to deposits of soot from the burning of its tungsten filament during the production of light. High-intensity discharge (HID) lamps, like the ones found in projection systems, do more with less by using a smaller electrical discharge to produce far more visible light. Expect to pay considerably more for projector bulbs than for standard bulbs of a comparable wattage.

Real World Scenario

Burning Up

Although it doesn't take long for the fan to stop running and the projector to cool down, this phase should never be skipped to save time. With projector bulbs being one of the priciest consumables in the world of technology, doing so may cost you more than a change in your travel arrangements.

A fellow instructor relayed the story of carrying his own portable projector with him on the road. At the end of a week's class, he would power down the projector and get his laptop

and other goodies packed away. Just before running out of the door, he would unplug the projector and pack it up. As with many instructors, this gentleman's presentations increased in density and length as he became more and more comfortable with the material.

Over time, his presentation had been running later and later each Friday afternoon, edging him ever closer to his airline departure time. He admitted that he had gotten into the habit of yanking the power plug for his projector from the wall and quickly stuffing the unit into the carrying case before darting out the door. It's no wonder then that a few months after he started this behavior, his projector failed catastrophically. Replacing the bulb was not the solution.

One caveat with projectors is that you must never pull the electrical plug from the outlet until you hear the internal fan cut off. There is enough residual heat generated by the projector bulb that damage to the electronics or the bulb itself (discoloration or outright failure) can occur if the fan is not allowed to remove enough heat before it stops running. Without a connection to an electrical outlet, the fan stops immediately. The electronics have the appropriate level of heat shielding so that the fan removes enough heat during normal operation to avoid damage to the shielded components.

Audio Devices

Audio devices, true to their name, produce sound by plugging into a sound card. Many sound cards today are integrated into a device's motherboard, but some computers still have separate audio expansion cards. Audio devices can provide output, such as through speakers or headphones, or input with a microphone.

Speakers and headphones generally connect with a 1/8″ (3.5mm) audio connector, as shown in Figure 3.6. You may recall seeing the ports for these in Volume 1 Chapter 2, "Expansion Cards, Storage Devices, and Power Supplies," when we covered sound cards. Most audio connectors have two thin black bands engraved on them, which separates the connector into three parts: the tip, ring, and sleeve. Because of this, sometimes you will see these connectors referred to as *TRS connectors*. The tip provides left audio, the first band above the black groove (the ring) provides right audio, and the sleeve is the ground. You'll notice that the audio connector in Figure 3.6 has three black bands, providing four connections and making it a TRRS connector. The fourth one is for the microphone.

FIGURE 3.6 1/8″ audio connector

Headsets that provide audio and a microphone are popular for audio conferencing calls and video gaming. A sample headset is shown in Figure 3.7. This model connects via USB, as do most headsets. Volume controls and a microphone mute are located on the right earpiece.

FIGURE 3.7 A USB headset

Although discussed throughout this chapter, the *microphone* has yet to be formally defined, a definition that is at once technical and simple. Microphones convert sound waves into varying electrical signals. The result can be recorded, transmitted, or altered in a variety of ways, including amplification.

When installing a microphone, you must match its connector with an available one on the computer. Modern choices include the classic analog pink TRS connector and USB. Wireless versions also exist, but their receiver might still be connected to a standard I/O port. Alternatively, the microphone could be paired with a built-in Bluetooth transceiver, headphones, or headset.

Configuring a microphone on a PC can be done in a few places in Windows, depending on what you need to do. For microphone configuration options such as input sensitivity, look for the Recording tab of the Sound applet in Control Panel. Options here include setting the levels and choosing enhancements, such as noise suppression and echo cancellation. To select which apps can use the microphone, go to Start ➢ Settings ➢ Privacy & Security ➢ Microphone. Specialized applications may also have internal configuration for the microphone, passing most details of the configuration back to the operating system.

Input and Output Devices

An *input device* is one that transfers information from outside the computer system to an internal storage location, such as system RAM, video RAM, flash memory, or disk storage. Without input devices, computers would be unable to change from their default boot-up state. An *output device* does the opposite of an input device—it takes information that's stored in RAM or another location and spits it back out somewhere for the user to do something with it. We've already covered monitors, which are the most common output device. The other major type of output device is a printer. Volume 1 Chapter 5, "Printers and Multifunction Devices," is dedicated to them. Further, some devices are capable of managing both input and output.

Frequently, when discussed together, input and output are abbreviated as I/O.

Keyboard

The *keyboard* is easily the most popular PC input device, so much so that it's more of a necessity. Very few users would even think of beginning a computing session without a working keyboard. Fewer still would even know how. The U.S. English keyboard places keys in the same orientation as the QWERTY typewriter keyboards, which were developed in the 1860s. Wired keyboards are almost always attached via USB. Wireless keyboards will often have a USB dongle that is attached to the computer, but they can also use Bluetooth.

Keyboards have also added separate number pads to the side and function keys (not to be confused with the common laptop key labeled *Fn*), placed in a row across the top of the keyboard above the numerical row. Key functionality can be modified by using one or more combinations of the Ctrl, Alt, Shift, and laptop Fn keys along with the normal QWERTY keys.

In addition to the layout for a standard keyboard, other keyboard layouts exist—some not nearly as popular. For example, without changing the order of the keys, an ergonomic keyboard is designed to feel more comfortable to users as they type. The typical human's hands do not rest with the fingers straight down. Ergonomic keyboards, therefore, should not place keys flat and along the same plane. To accomplish that goal, manufacturers split the keyboard down the middle, angling keys on each side downward from the center. Doing so fits the keys to the fingers of the hands when they are in a relaxed state. Figure 3.8 shows an example of an ergonomic keyboard. Even more exotic-looking ergonomic keyboards exist and may provide better relief for users who suffer from repetitive-use problems such as carpal tunnel.

Here are two places to go for keyboard-related settings in Windows 11:

- Keyboard properties, by typing **keyboard** into the Windows search box and pressing Enter
- Typing settings, by clicking Start ➢ Settings ➢ Time & Language ➢ Typing

FIGURE 3.8 An ergonomic keyboard

Mouse

Although the computer *mouse* was born in the 1970s at Xerox's Palo Alto Research Center (PARC), it was in 1984 that Apple made the mouse an integral part of the personal computer with the introduction of the Macintosh. In its most basic form, the mouse is a hand-fitting device that uses some form of motion-detection mechanism to translate its own physical two-dimensional movement into onscreen cursor motion. Many variations of the mouse exist, including trackballs, tablets, touchpads, and pointing sticks. Figure 3.9 illustrates a basic wireless optical mouse.

FIGURE 3.9 A computer mouse

The motion-detection mechanism of the original Apple mouse was a simple ball that protruded from the bottom of the device so that when the bottom was placed against a flat surface that offered a slight amount of friction, the mouse would glide over the surface but the ball would roll, actuating two rollers that mapped the linear movement to a Cartesian plane and transmitted the results to the software interface. This method of motion detection has been replaced by optical receptors to catch LED light reflected from the surface the mouse is used on. Note that most optical mice will have problems working on a transparent glass surface because of the lack of reflectivity.

The mouse today can be wired to the computer system or connected wirelessly. A wired mouse typically uses a USB port, which also provides power. Wireless versions will have a USB dongle or connect via Bluetooth. They are powered with batteries, and the optical varieties deplete these batteries more quickly than their mechanical counterparts.

The final topic is one that is relevant for any mouse: buttons. The number of buttons that you need your mouse to have depends on the software interfaces you use. For the Macintosh, one button has always been sufficient, but for a Windows-based computer, at least two are recommended—hence, the term *right-click*. Today, the mouse is commonly found to have a wheel on top to aid in scrolling and other specialty movement. The wheel has even developed a click in many models, sort of an additional button underneath the wheel. Buttons on the side of the mouse that can be programmed for whatever the user desires are common today as well.

There are several variants on pointer devices, such as trackballs. A *trackball* is like an inverted mouse. Both devices place the buttons on the top, which is where your fingers will be. A mouse places its tracking mechanism on the bottom, requiring that you move the entire assembly as an analogue for how you want the cursor on the screen to move. In contrast, a trackball places the tracking mechanism, usually a ball that is about one inch in diameter, on the top with the buttons. You then have a device that need not be moved around on the desktop and can work in tight spaces and on surfaces that would be incompatible with the use of a mouse. The better trackballs place the ball and buttons in such a configuration that your hand rests ergonomically on the device, allowing effortless control of the onscreen cursor.

Other input devices you may encounter are touchpads (especially on laptops), signature pads, game controllers, barcode and QR code scanners, magnetic or chip readers, and near-field communication (NFC) devices such as tap-to-pay devices. They almost always connect to the computer using a USB port.

KVM Switches

Some users will need multiple computers in the same location, but they don't want to hassle with multiple sets of input and output devices. For them, a KVM switch may be a good solution.

A KVM switch isn't an input or output device per se, but it allows you to switch between sets of input and output devices. The *KVM switch* is named after the devices among which it allows you to switch. The initials stand for *keyboard*, *video*, and *mouse*. KVM switches come in a variety of models, with different connector types available.

The purpose of the switch is to allow you to have multiple computers attached to the same keyboard, monitor, and mouse. You can use these three devices with only one system at a time. Some switches have a dial that you turn to select which system attaches to the components, while others feature buttons for each system connected. Common uses of KVM switches include using the same components alternately for a desktop computer and a laptop docking station or having a server room with multiple servers but no need to interface with them simultaneously.

Storage Devices

We spent quite a lot of time in Volume 1 Chapter 2 discussing storage options, such as hard drives and optical drives. These devices are frequently internal to the case, but external options are available as well.

Take optical drives, for instance. In order to save space and reduce costs, manufacturers usually don't include internal optical drives. If users want to play a Blu-ray or DVD movie, they will need to attach an external optical drive. External optical drives can be used for data backups as well. These external drives will most likely connect via USB or eSATA.

External storage drives can greatly enhance the storage capacity of a computer, or they can provide networked storage for several users. A plethora of options is available, from single drives to multidrive systems with several terabytes of capacity. Figure 3.10 shows an external network-attached storage (NAS) device.

Looking at Figure 3.10, you can see that this is a self-enclosed unit that can hold up to five hard drives, which are externally accessible—a nice feature. Some hold more; some hold fewer. This NAS also has status lights next to the drives to indicate proper operation or failure as well as a small status display. Nicer NAS systems enable you to hot-swap hard drives and have built-in fault tolerance as well.

In addition to the hardware, the NAS device contains its own operating system, meaning that it acts like its own file server. In most cases, you can plug it in, do some very minor configuration, and have instant storage space on your network. As far as connectivity goes, NAS systems will connect to a PC through a USB or eSATA port, but that is primarily so you can use that PC to run the configuration software for the NAS. The NAS also connects to the network, and that is how all network users access the storage space.

Understanding Cables and Connectors

Peripheral devices used with a computer need to attach to the motherboard somehow. They do so through the use of ports and cables. A *port* is a generic name for any connector on a computer or peripheral into which a cable can be plugged. A *cable* is simply a way of

connecting a peripheral or other device to a computer using multiple copper or fiber-optic conductors inside a common wrapping or sheath. Typically, cables connect two ports: one on the computer and one on some other device.

FIGURE 3.10 A network-attached storage device

Source: Bin im Garten / Wikimedia Commons / CC BY SA 3.0

The A+ exam objectives break cables and connectors into two different subobjectives, but really they need to be discussed together. After all, a cable without a connector doesn't do much good, and neither does a connector without a cable. In the following sections, we'll look at three different classifications of cables and the connectors that go with them: peripheral, video, and hard drive.

For the A+ exam, you will also need to be familiar with network cables and connectors. We will cover those in depth in Volume 1 Chapter 6, "Networking Fundamentals."

Peripheral Cables and Connectors

Some cables are for specific types of devices only. For example, SATA is for storage devices. And HDMI can transmit audio as well as video, but for the most part, though, we associate HDMI with video.

Unlike HDMI and SATA, the cables and connectors in this section are specifically designed to connect a variety of devices. For example, someone may have a USB hub with a wireless mouse, network card, Lightning cable (to charge an iPhone), and flash drive all attached at the same time. Those four devices serve very different purposes, but they all share the USB connection in common. We'll start with the highly popular USB and then discuss Lightning ports, Thunderbolt cables, and serial cables.

Exam Objective 3.2, "Summarize basic cable types and their connectors, features, and purposes," lists four peripheral cable types you should know: USB 2.0, USB 3.0, Serial, and Thunderbolt.

Universal Serial Bus

It seems like we've mentioned USB a thousand times already in this book, because it's everywhere and it's like everyone knows what it is. But we have yet to define or explain it. *Universal Serial Bus (USB)* is a standard used to connect a wide variety of peripherals, such as keyboards, mice, digital cameras, printers, scanners, hard drives, and network cards, to computers. USB was designed by several companies, including Intel, Microsoft, and IBM, and is currently maintained by the USB Implementers Forum (USB-IF).

USB technology is fairly straightforward. Essentially, it is designed to be Plug and Play—just plug in the peripheral and it should work, provided that the software is installed to support it. Many standard devices have drivers that are built into the common operating systems or automatically downloaded during installation. More complex devices come with drivers to be installed before the component is connected.

USB host controllers can support up to 127 devices, which is accomplished through the use of a 7-bit identifier. (A bit has two values, 0 and 1. Therefore, a string of 7 consecutive bits has 2^7 possible unique values, or 128 total.) The 128th identifier, the highest address, is used for broadcasting to all endpoints. Realistically speaking, you'll probably never get close to this maximum. Even if you wanted to try, you won't find any computers with 127 USB ports. Instead, you would plug in a device known as a *USB hub* (shown in Figure 3.11) into one of your computer's USB ports, which will give you several more USB ports from one original port. Understand that a hub counts as a device for addressing purposes. Hubs can be connected to each other, but interconnection of host controllers is not allowed; each one and its connected devices are isolated from other host controllers and their devices. As a result, USB ports are not considered networkable ports. Consult your system's documentation to find out if your USB ports operate on the same host controller.

Another nice feature of USB is that devices can draw their power from the USB cable, so you may not need to plug in a separate power cord. This isn't universally true, though, as some peripherals still require external power.

USB Standards

Even though USB was released in 1996, the first widely used standard was USB 1.1, which was released in 1998. It was pretty slow—only 12 Mbps at full speed and 1.5 Mbps at low speed—so it was only used for keyboards, mice, and printers. When USB 2.0 came out in 2000 with a faster transfer rate of 480 MBps (called *Hi-Speed*), video devices were possible.

The newer USB 3.x, USB4, and USB4 v2 standards have increased throughput even further. Table 3.1 lays out the specifications and speeds for you.

FIGURE 3.11 A 4-port USB hub

TABLE 3.1 USB specifications

Specification	Release year	Maximum speed	Trade name	Color
USB 1.1	1998	12 Mbps	Full-Speed	White
USB 2.0	2000	480 Mbps	Hi-Speed	Black
USB 3.0	2008	5 Gbps	SuperSpeed	Blue
USB 3.1	2013	10 Gbps	SuperSpeed+	Teal
USB 3.2	2017	20 Gbps	SuperSpeed+	Red
USB4	2019	40 Gbps	USB4 40 Gbps	n/a
USB4 2.0	2022	80 Gbps	USB4 80 Gbps or USB4 v2	n/a

The USB 1.x and 2.x specifications didn't recommend a specific color for the ports, but when USB 3.0 was released, the USB Implementers Forum suggested that the ports and cable connectors be colored blue, to signify that they were capable of handling higher speeds. Device manufacturers are not required to follow the color-coding scheme, so you may see some inconsistency. For example, if you look at image 3.12, you will see 10 USB ports on the back of an ASUS motherboard. We know the book prints in black-and-white, so bear with us as we describe the scene. The top row of 4 ports are all black, so you can assume these

are USB 2.0. Below the HDMI port, there are two USB ports, white and blue, and both are labeled USB 3.2. The third row of USB ports are also labeled 3.2 and are red, and the fourth row are also labeled USB 3.2 but are blue. One can assume in this case that the bottom six USB ports all support USB 3.2, regardless of their color.

FIGURE 3.12 USB ports on a motherboard

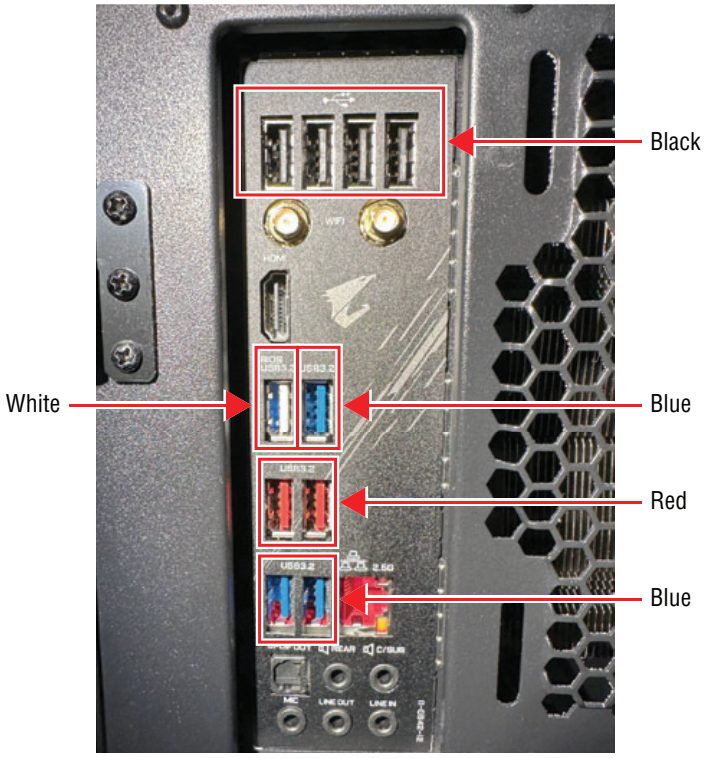

A yellow USB port is "always on," meaning it's capable of charging a connected device even if the PC is sleeping or shut down. USB4 v2 is the newest standard, and in many ways is comparable to Thunderbolt 5 (discussed later). USB4 v2 is typically implemented over USB-C cables.

USB Power

As mentioned previously, USB ports provide power to devices plugged into them. Typical power for attached USB devices is 5V. The maximum current (amps) and wattage will depend on the connected device and USB standard being used.

All USB ports are also capable of functioning as charging ports for devices such as tablets, smartphones, and smart watches. The charging standard, called *USB Battery Charging*, was released in 2007. USB Power Delivery (PD) was developed in 2012. Technically, they are different standards, but in practice, USB ports are capable of supporting both standards at the same time. Table 3.2 outlines some of the versions and the maximum power that they provide. The newest version, USB PD 3.1, requires the use of a USB-C cable.

TABLE 3.2 USB power standards

Standard	Year	Maximum power
USB Battery Charging 1.0	2007	5V, 1.5A (7.5W)
USB Battery Charging 1.2	2010	5V, 5A (20W)
USB Power Delivery 1.0	2012	20V, 5A (100W)
USB Power Delivery 2.0 (specified use of Type-C connectors but only up to 15W)	2014	5V, 3A (15W)
USB Power Delivery 3.0	2015	20V, 5A (100W)
USB Power Delivery 3.1	2021	48V, 5A (240W)

A smartphone or tablet typically needs a minimum of about 7.5 watts to charge properly. The Battery Charging 1.0 standard was good enough, but not for larger devices. For example, about 20 watts is required to power a small laptop computer, and standard 15-inch laptops can require 60 watts or more. With USB PD, one USB port can now provide enough power for a laptop as well as a small printer.

Because of the capabilities of USB PD, it's becoming common to see devices up to laptop size lose their standard AC power ports and adapters—they may just have a USB-C port instead. To get the full capabilities of USB PD, you need to use a USB-C port and cable.

USB Cables and Connectors

In order to achieve the full speed of the specification that a device supports, the USB cable needs to meet that specification as well. In other words, USB 1.x cables cannot provide USB 2.0 and 3.x performance, and USB 2.0 cables cannot provide USB 3.x performance. Otherwise, the connected device will have to fall back to the maximum version supported by the cable. This is usually not an issue, except for the lost performance, but some high-performance devices will refuse to operate at reduced levels. Note that all specifications are capable of Low Speed, which is a 1.5 Mbps performance standard that has existed since the beginning of USB time.

Throughout most of its history, USB has relied on a small suite of standard connectors. The two broad classifications of connectors are designated Type-A and Type-B connectors, and there are micro and mini versions of each. A standard USB cable has some form of Type-A connector on the end that plugs into the computer or hub, and some form of Type-B or proprietary connector on the device end. Figure 3.13 shows four classic USB 1.x/2.0 cable connectors. From left to right, they are as follows:

- Micro-USB
- Mini-USB
- Type-B
- Type-A male

FIGURE 3.13 Standard USB connectors

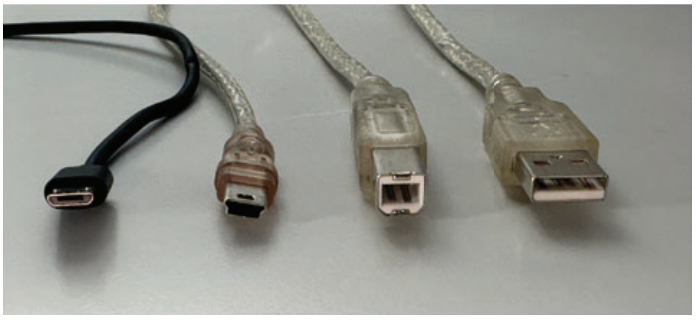

Small form factor devices, including many smartphones and smaller digital cameras, use a *micro-USB* or *mini-USB* connector, unless the manufacturer has developed its own proprietary connector. Micro-USB connectors (and modified ones) are popular with many Android phone manufacturers.

In 2014, a new connector named *USB Type-C* (or simply *USB-C*) was developed. USB-C is designed to replace Type-A and Type-B, and, unlike its predecessors, it's reversible. That means no more flipping the connector over several times to figure out which way it connects. Type-C cables will also be able to provide more power to devices than classic cables were. Figure 3.14 shows a Type-C connector and a Type-A connector. You can see that while the Type-A connector is rectangular-shaped, the Type-C connector has rounded corners and looks more like an elongated oval.

FIGURE 3.14 USB Type-C (top) and Type-A (bottom)

One point of confusion for many is dissociating the connector type from the standard. Because USB 3.1 and USB-C were both released around the same time, people often think that they are one and the same—but they're not. USB 3.1 can be implemented using classic A and B connectors, and USB 2.0 can work over a Type-C connector.

USB was designed to be a short-distance technology. Because of this, USB cables are limited in length. USB 1.x and 2.0 can use cables up to 5 meters long, whereas USB 3.x can use cables up to 3 meters long. The maximum length of a USB4 cable is even shorter yet at 80 centimeters (0.8 meters). In addition, if you use hubs, you should never use more than five hubs between the system and any component.

Despite the seemingly locked-up logic of USB connectivity, it is occasionally necessary to alter the interface type at one end of a USB cable. For that reason, there are a variety of simple, passive converters on the market with a USB interface on one side and a USB or different interface on the other. Along with adapters that convert USB Type-A to USB Type-B, there are adapters that will convert a male connector to a female one. In addition, you can convert USB to a lot of other connector types, such as USB to Ethernet (shown in Figure 3.15), USB to SATA, USB to eSATA, USB to PS/2, USB to serial, and a variety of others.

FIGURE 3.15 Kensington USB-to-Ethernet adapter

For more information on USB, check out www.usb.org.

Lightning

Introduced in 2012 with the iPhone 5, the *Lightning connector* was Apple's proprietary connector for iPads until the third-generation iPad Pro in 2018 and iPhones up until the iPhone 15. It's an 8-pin connector that replaced Apple's previous 30-pin dock connector. A standard Lightning cable has a USB Type-A or USB-C connector on one end and the Lightning connector on the other, as shown in Figure 3.16. It's not keyed, meaning that you can put it in with either edge up.

FIGURE 3.16 Lightning cable

Lightning cables support USB 2.0. You will find cables that are USB-C to Lightning, as well as various Lightning adapters, such as those to HDMI, DisplayPort, audio, and Lightning to female USB Type-A (so you can plug a USB device into an iPad or iPhone).

Thunderbolt

Where there's lightning, there's thunder, right? Bad joke attempts aside, in computer circles Lightning connectors don't have anything to do with Thunder(bolt). *Thunderbolt*, created in collaboration between Intel and Apple and released in 2011, combines PCI Express 2.0 x4 with the DisplayPort 1.x technology. While it's primarily used for video (to replace DisplayPort), the connection itself can support multiple types of peripherals, much like USB does.

Thunderbolt Standards

For most of their histories, Thunderbolt and USB have been competing standards. Thunderbolt was designed more for video applications and USB was the slower "jack of all trades" port, but in reality, they could be used for almost the exact same list of peripherals. It just depended on what your computer supported. But as we pointed out in the USB section, the new USB4 version is based on Thunderbolt 3, providing the same speed and using the same connectors. Table 3.3 shows the five Thunderbolt versions and some key characteristics.

TABLE 3.3 Thunderbolt standards

Version	Year	Maximum throughput	Connector	Other new features
Thunderbolt 1	2011	10 Gbps	Mini DisplayPort	
Thunderbolt 2	2013	20 Gbps	Mini DisplayPort	DisplayPort 1.2 (can send video to a 4k display)
Thunderbolt 3	2015	40 Gbps	USB-C	10 Gbps Ethernet support
Thunderbolt 4	2020	40 Gbps	USB-C	Can support two 4k displays or one 8k display; 32 Gbps PCIe
Thunderbolt 5	2024	80/120 Gbps	USB-C	Can support multiple 8k displays and 540 Hz refresh rates, or 3 displays at 144 Hz

Thunderbolt 3 was released in 2015 and doubled the bandwidth to 40 Gbps. It supports PCIe 3.0 and DisplayPort 1.2, meaning that it can support dual 4K displays at 60 Hz or a single 4K display at 120 Hz. It also provides up to 100 watts of power to a device.

Thunderbolt 4 was released in 2020. Perhaps the most interesting thing about the release is what it didn't do, which is increase data transfer rates versus Thunderbolt 3. It still had a maximum bandwidth of 40 Gbps. And the maximum of 100 watts of power to attached devices didn't change either. The big advantages of Thunderbolt 4 included support for two 4k displays or one 8k display and the requirement to support 32 Gbps data transfers via PCIe, up from 16 Gbps in version 3. Version 5, released in 2024, doubles the speed of version 4.0 to 80 Gbps, and can be boosted to 120 Gbps for limited amounts of time. It also supports multiple 8k displays at once, and it can provide 240 watts of charging power.

Thunderbolt Cables and Connectors

The most common Thunderbolt cable is a copper, powered active cable extending as far as 3 meters, which was designed to be less expensive than an active version of a DisplayPort cable of the same length. There are also optical cables in the specification that can reach as far as 60 meters. Copper cables can provide power to attached devices, but optical cables can't.

Additionally, and as is the case with USB, Thunderbolt devices can be daisy-chained and connected via hubs. Daisy chains can extend six levels deep for each controller interface, and each interface can optionally drive a separate monitor, which should be placed alone on the controller's interface or at the end of a chain of components attached to the interface.

As noted in Table 3.3, Thunderbolt changed connectors between versions 2 and 3. Figure 3.17 shows two Thunderbolt 2 interfaces next to a USB port on an Apple MacBook Pro. Note the standard lightning-bolt insignia by the port. Despite its diminutive size, the Thunderbolt port has 20 pins around its connector bar, like its larger DisplayPort cousin. Of course, the functions of all the pins do not directly correspond between the two interface types, because Thunderbolt adds PCIe functionality.

FIGURE 3.17 Two Thunderbolt 2 interfaces

Starting with Thunderbolt 3, the connector was changed to standard USB-C connectors, as shown in Figure 3.18. Notice that the lightning bolt icon remained the same.

Converters are available that connect Thunderbolt connectors to VGA, HDMI, and DVI monitors. Active converters that contain chips to perform the conversion are necessary in situations such as when the technology is not directly pin-compatible with Thunderbolt—as with VGA and DVI-A analog monitor inputs, for example. Active converters are only slightly more expensive than their passive counterparts but still only a fraction of the cost of Thunderbolt hubs. One other advantage of active connectors is that they can support resolutions of 4k (3840 × 2160) and higher.

FIGURE 3.18 Two Thunderbolt 3 interfaces

Source: Amin / Wikimedia Commons / CC BY SA 4.0

Serial Ports

Before USB came along in 1998, *serial ports* were considered slow and inferior to parallel ports. Still, serial enjoyed use among peripherals that didn't need to transfer information at high speeds, such as mice, modems, network management devices, and even printers. Figure 3.19 shows a 9-pin serial port. It's the one marked "Serial," and it's also the only male connector on the back of the PC.

FIGURE 3.19 Several peripheral ports

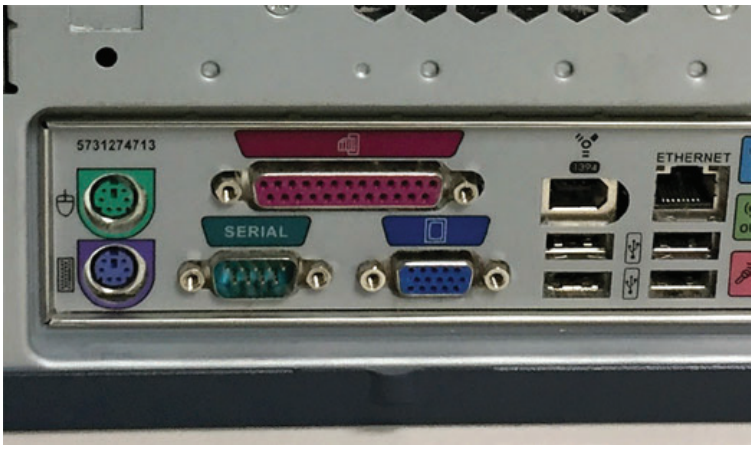

As you might expect, a serial cable attaches to the serial port. Figure 3.20 shows a female DB-9 serial connector. To make things more confusing, sometimes you will hear people refer to the image in Figure 3.20 as an RS-232 cable or connector. Even though the terms are often used interchangeably, there is a technical difference.

FIGURE 3.20 DB-9 serial connector

DB-9 refers to a specific type of D-sub connector that has 9 pins. *RS-232*, on the other hand, is a communications standard for serial transmission. In other words, systems may communicate with each other using RS-232 over a DB-9 connection. But RS-232 can be used on other types of serial cables as well, such as DB-15 or DB-25. Generally speaking, if someone asks for an RS-232 serial cable, they *mean* a DB-9 cable with female connectors. But it's always best to confirm.

RS-232 did have a few advantages over USB—namely, longer cable length (15 meters vs. 3–5 meters) and a better resistance to electromagnetic interference (EMI). Still, USB has made old-school serial ports nearly obsolete. About the only time they are used today is for management devices that connect to servers or network routers with no keyboard and monitor installed.

Other Keyboard and Mouse Connectors

At one time, the most popular connectors for keyboards and mice were round connectors called *Personal System/2 (PS/2)* connectors. The PS/2 connector (there are two on the left in Figure 3.19) is a smaller 6-pin mini-DIN connector. Many PCs included a PS/2 keyboard connector as well as a PS/2 mouse connector right above it on the motherboard. The keyboard connector was colored purple, and the mouse one green. The ends of the keyboard and mouse cables would be purple and green as well. Today, the PS/2-style connector has been replaced by the USB port.

If you do run into a keyboard or mouse with a PS/2 connector, or a motherboard with a PS/2 port, you can buy a PS/2-to-USB adapter to make them work with more current hardware.

Video Cables and Connectors

Computer displays are ubiquitous—they're easily the most widely used peripheral. Different standards exist to connect displays to the computer, and you need to be familiar with five of them for the exam. We will start with the older technologies and work toward the present.

Exam Objective 3.2 lists VGA, DVI, HDMI, DisplayPort, and USB-C as video cables and connectors to know.

Video Graphics Array Connector

The *Video Graphics Array (VGA) connector* was the de facto video standard for computers for years and is still in use today. First introduced in 1987 by IBM, it was quickly adopted by other PC manufacturers. The term *VGA* is often used interchangeably to refer to generic analog video, the 15-pin video connector, or a 640 × 480 screen resolution (even though the VGA standard can support much higher resolutions). Figure 3.21 shows a VGA port, as well as the male connector that plugs into the port. Nearly all VGA connectors are blue.

FIGURE 3.21 VGA connector and port

Understanding D-Sub Ports and Connectors

The VGA connector is an example of a *D-subminiature connector*, also known as a *D-sub connector*. For a number of years, D-sub was the most common style of connector found on computers. Their names are typically designated with D*X-n*, where the letter *X* is replaced by a letter from *A* to *E*, which refers to the size of the connector, and the letter *n* is replaced by the number of pins or sockets in the connector.

D-sub connectors are usually shaped like a trapezoid and have at least two rows of pins with no other keying structure or landmark. Several were shown back in Figure 3.19. At the bottom center in Figure 3.19 is a DE-15F 15-pin display-connector port, which may also be referred to as an HD-15 or DB-15 port. The top one is a classic DB-25 parallel port, and the

bottom left is a *DB-9* serial port. The IEEE 1394 (FireWire—now obsolete) and two dusty USB ports are shown for a size comparison.

The "D" shape ensures that only one orientation is possible. If you try to connect them upside down or try to connect a male connector to another male connector, they just won't go together and the connection can't be made. By the way, male interfaces have pins, while female interfaces have sockets. Be on the lookout for the casual use of *DB* to represent any D-sub connector. This is very common and is accepted as an unwritten de facto standard, even if some are technically DE- or DA- connectors. Also note that you will see them written without the hyphen or with a space, such as DB15 or DB 15.

VGA technology is the only one on the objectives list that is purely analog. It has been superseded by newer digital standards, such as DVI, HDMI, and DisplayPort, and it was supposed to be phased out starting in 2013. A technology this widely used will be around for quite a while, though, and you'll still see it occasionally in the wild (or still in use).

All the video connector types introduced from here on are digital standards.

Digital Visual Interface

The analog VGA standard ruled the roost for well over a decade but it had a lot of shortcomings. Digital video can be transmitted farther and at higher quality than analog, so development of digital video standards kicked off in earnest. The first commercially available one was a series of connectors known collectively as *Digital Visual Interface (DVI)* and was released in 1999.

At first glance, the DVI connector might look like a standard D-sub connector. On closer inspection, however, it begins to look somewhat different. For one thing, it has quite a few pins, and for another, the pins it has are asymmetrical in their placement on the connector. The DVI connector is usually white and about an inch long. Figure 3.22 shows what the connector looks like coming from the monitor.

FIGURE 3.22 DVI connector

Source: Evan Amos / Wikimedia Commons / Public domain

There are three main categories of DVI connectors:

DVI-A *DVI-A* is an analog-only connector. The source must produce analog output, and the monitor must understand analog input.

DVI-D *DVI-D* is a digital-only connector. The source must produce digital output, and the monitor must understand digital input.

DVI-I *DVI-I* is a combination analog/digital connector. The source and monitor must both support the same technology, but this cable works with either a digital or an analog signal.

The DVI-D and DVI-I connectors come in two varieties: single-link and dual-link. The dual-link options have more conductors—taking into account the six center conductors—than their single-link counterparts; therefore, the dual-link connectors accommodate higher speed and signal quality. The additional link can be used to increase screen resolution for devices that support it. Figure 3.23 illustrates the five types of connectors that the DVI standard specifies.

FIGURE 3.23 Types of DVI connectors

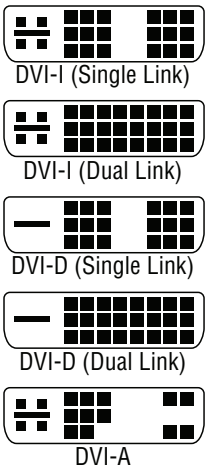

DVI-A and DVI-I analog quality is superior to that of VGA, but it's still analog, meaning that it is more susceptible to noise. However, the DVI analog signal will travel farther than the VGA signal before degrading beyond usability. Nevertheless, the DVI-A and VGA interfaces are pin-compatible, meaning that a simple passive *DVI-to-VGA adapter*, as shown in Figure 3.24, is all that is necessary to convert between the two. As you can see, the analog portion of the connector, if it exists, comprises the four separate color and sync pins and the horizontal blade that they surround, which happens to be the analog ground lead that acts as a ground and physical support mechanism even for DVI-D connectors.

FIGURE 3.24 DVI-to-VGA adapter

It's important to note that DVI-I cables and interfaces are designed to interconnect two analog or two digital devices; they cannot convert between analog and digital. DVI cables must support a signal of at least 4.5 meters, but better cable assemblies, stronger transmitters, and active boosters result in signals extending over longer distances.

One thing to note about analog versus digital display technologies is that all graphics' adapters and all monitors deal with digital information. It is only the connectors and cabling that can be made to support analog transmission. Before DVI and HDMI encoding technologies were developed, consumer digital video display connectors could not afford the space to accommodate the number of pins that would have been required to transmit 16 or more bits of color information per pixel. For this reason, the relatively few conductors of the inferior analog signaling in VGA were appealing.

High-Definition Multimedia Interface

High-Definition Multimedia Interface (HDMI) is an all-digital technology that advances the work of DVI to include the same dual-link resolutions using a standard HDMI cable but with higher motion-picture frame rates and digital audio right on the same connector. HDMI was introduced in 2002, which makes it seem kind of old in technology years, but it's a great, fast, reliable connector that will probably be around for several years to come. HDMI cabling also supports an optional Consumer Electronics Control (CEC) feature that allows transmission of signals from a remote-control unit to control multiple devices without separate cabling to carry infrared signals.

HDMI cables, known as Standard and High Speed, exist today in the consumer space. Standard cables are rated for 720p resolution as well as 1080i, but not 1080p. High Speed cables are capable of supporting not only 1080p, but also the newer 4k and 8k technologies. Figure 3.25 shows an HDMI port and cable.

In June 2006, revision 1.3 of the HDMI specification was released to support the bit rates necessary for HD DVD and Blu-ray Disc. This version also introduced support for "deep color," or color depths of at least one billion colors, including 30-, 36-, and 48-bit color. However, not until version 1.4, which was released in May 2009, was the High Speed HDMI cable initially required.

With version 1.4 came HDMI capability for the controlling system—the television, for instance—to relay Ethernet frames between its connected components and the Internet, alleviating the need for each and every component to find its own access to the LAN for Internet

access. Both Standard and High Speed cables are available with this Ethernet channel. Each device connected by such a cable must also support the HDMI Ethernet Channel specification, however.

FIGURE 3.25 HDMI port and cable

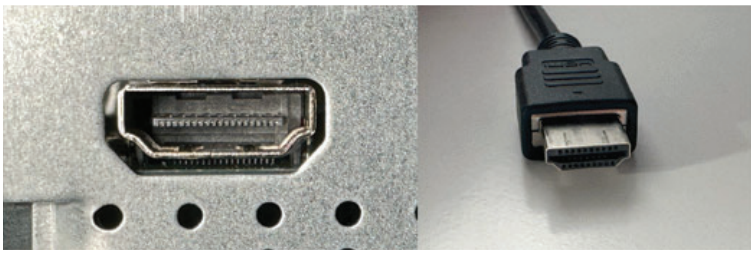

Additional advances that were first seen in version 1.4 were 3D support, 4K resolution (but only at a 30 Hz refresh rate), an increased 120 Hz refresh rate for the 1080 resolutions, and an Audio Return Channel (ARC) for televisions with built-in tuners to send audio back to an A/V receiver without using a separate output cable. Version 1.4 also introduced the anti-vibration Type-E locking connector for the automotive-video industry and cables that can also withstand vibration as well as the hot/cold extremes that are common in the automotive world.

Version 2.0 of HDMI (2013) introduced no new cable requirements. In other words, the existing High Speed HDMI cable is fully capable of supporting all new version 2 enhancements. These enhancements include increasing the 4K refresh rate to 60 Hz, a 21:9 theatrical widescreen aspect ratio, and 32-channel audio. Note that 7.1 surround sound comprises only eight channels, supporting the more lifelike Rec. 2020 color space and multiple video and audio streams to the same output device for multiple users. Version 2.0a, released in 2015, primarily added high dynamic range (HDR) video, but it does not require any new cables or connectors.

The most recent version (as of this writing) is HDMI 2.1, released in November 2017. Version 2.1 specifies a new cable type called 48G, which provides for 48 Gbps bandwidth. 48G cables are backward compatible with older HDMI versions. You can also use older cables with 48G-capable devices, but you just won't get the full 48 Gbps bandwidth. HDMI 2.1 also provides for 120 Hz refresh rates for 4k, 8k, and 10k video, and it supports enhanced Audio Return Channel (eARC), which is needed for object-based audio formats, such as DTS:X and Dolby Atmos.

Even though the HDMI connector is not the same as the one used for DVI, the two technologies are electrically compatible. HDMI is compatible with DVI-D and DVI-I interfaces through proper adapters, but HDMI's audio and remote-control pass-through features are lost. Additionally, 3D video sources work only with HDMI. Figure 3.26 shows a *DVI-to-HDMI adapter* between DVI-D and the Type-A 19-pin HDMI interface. Compare the DVI-D interface in Figure 3.26 to the DVI-I interface in Figure 3.24, and note that the ground blade

on the DVI-D connector is narrower than that of the DVI-A and DVI-I connectors. The DVI-D receptacle does not accept the other two plugs for this reason; in addition, the four analog pins around the blade have no sockets in the DVI-D receptacle.

FIGURE 3.26 DVI-to-HDMI adapter

Unlike DVI-D—and, by extension DVI-I—DVI-A and VGA devices cannot be driven passively by HDMI ports directly. An *HDMI-to-VGA adapter* must be active in nature, powered either externally or through the HDMI interface itself.

HDMI cables should meet the signal requirements of the latest specification. As a result, and as with DVI, the maximum cable length is somewhat variable. For HDMI, cable length depends heavily on the materials used to construct the cable. Passive cables tend to extend no farther than 15 meters, while adding electronics within the cable to create an active version results in lengths as long as 30 meters.

Smaller HDMI connectors

Multiple versions of HDMI connectors are available in the marketplace. The standard connector that you're probably used to seeing, and the ones shown in Figure 3.25 and Figure 3.26, are the 19-pin Type-A connector. The Type-A connector and the 29-pin Type-B connector were specified in HDMI version 1.0 and haven't changed much since then. Type-B connectors were intended for higher-resolution products but are not used in the market today.

HDMI version 1.3 specified a smaller 19-pin Type-C connector for portable devices. The Type-C connector, also referred to as a *mini-HDMI connector*, is compatible with the Type-A connector, but it still requires an adapter due to its smaller size. HDMI version 1.4 specified two more interfaces: Type-D and Type-E. If Type-C is a mini-HDMI interface, then you might refer to the Type-D connector as micro-HDMI. Figure 3.27 shows all five HDMI connectors. Also compatible with Type-A interfaces because they have the same 19 pins, Type-D interfaces require just a simple adapter for conversion.

FIGURE 3.27 HDMI connector types

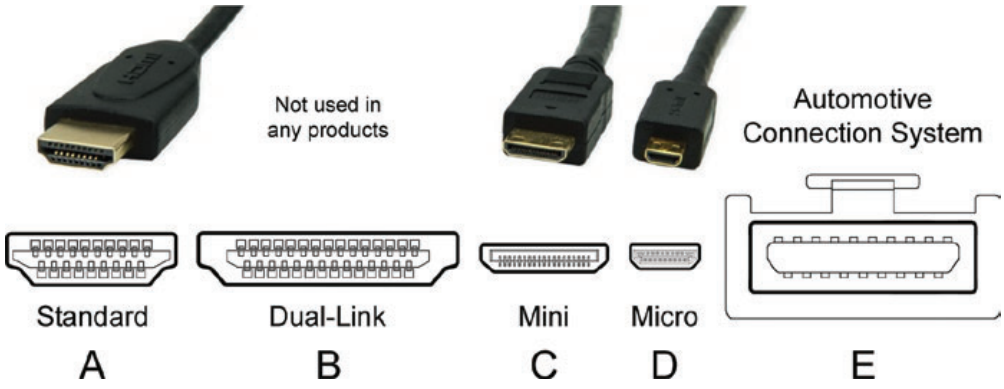

Source: C0nanPayne / Wikimedia Commons / Public domain

The mini-HDMI and micro-HDMI connectors are most often used on smaller portable devices, such as tablets, smartphones, and digital cameras. As mentioned previously, the Type-E connector has a locking mechanism and is intended for use in automobiles or other environments that are susceptible to vibration, which could cause a connector and cable to become disconnected.

DisplayPort

DisplayPort is a royalty-free digital display interface from the Video Electronics Standards Association (VESA) that uses less power than other digital interfaces and VGA. Introduced in 2008, it's designed to replace VGA and DVI. To help ease the transition, it's backward compatible with both standards, using an adapter. In addition, an adapter allows HDMI and DVI voltages to be lowered to those required by DisplayPort because it is functionally similar to HDMI and DVI. DisplayPort cables can extend 3 meters, unless an active cable powers the run, in which case the cable can extend to 33 meters. DisplayPort is intended primarily for video, but, like HDMI, it can transmit audio and video simultaneously.

Figure 3.28 shows a DisplayPort port as well as a connector. Note the beveled keying at the bottom-left corner of the port. The DisplayPort connector latches itself to the receptacle with two tiny hooks. A push-button mechanism serves to release the hooks for removal of the connector from the receptacle—you can see the button is on the top of the connector on the right in Figure 3.28.

The DisplayPort standard also specifies a smaller connector, known as the *Mini DisplayPort (mDP) connector*. The mDP is electrically equivalent to the full-size DP connector and features a beveled keying structure, but it lacks the latching mechanism present in the DP connector. Figure 3.29 shows an mDP to DisplayPort cable, so you can see both connectors at the same time.

FIGURE 3.28 A DisplayPort port and cable

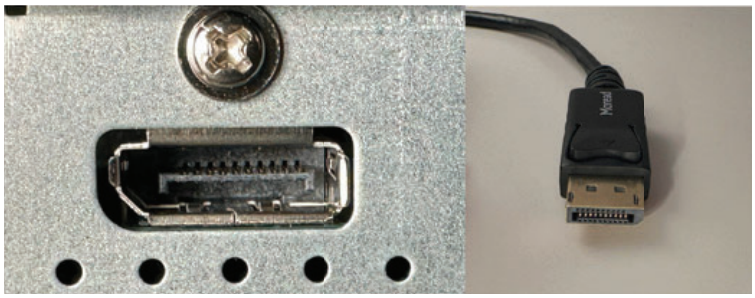

FIGURE 3.29 mDP and DisplayPort connectors

The final video cable connector to know for the exam is USB-C. Because we've already covered USB-C in depth earlier in this chapter, we won't repeat the information here. Do be aware that USB-C can be used for video connections as well as other peripherals.

Hard Drive Cables and Connectors

At the beginning of this chapter, we said that we were going to move outside the box and talk about external peripherals, cables, and connectors. For the most part that's true, but here we need to take a small digression to talk about connecting hard drives, most of which are internal. Some of this you already learned in Volume 1 Chapter 2, so this could feel like a review. Of course, there are SATA and PATA connectors, but we'll also throw in two new ones—SCSI and eSATA.

Optical drives use the same connectors as hard drives.

Remember that all drives need some form of connection to the motherboard so that the computer can "talk" to the disk drive. Regardless of whether the connection is built into the motherboard (*onboard*) or on an adapter card (*off-board*), internal or external, the standard for the attachment is based on the drive's requirements. These connections are known as *drive interfaces*. The interfaces consist of circuitry and a port, or *header*.

Serial Advanced Technology Attachment

The most common hard drive connector used today is *Serial Advanced Technology Attachment (SATA)*. Figure 3.30 shows SATA headers, which you have seen before, and a SATA cable. Note that the SATA cable is flat, and the connector is keyed to fit into the motherboard header in only one way. SATA data cables have a 7-pin connector. SATA power cables have 15 pins and are wider than the data connector.

FIGURE 3.30 SATA connectors and cable

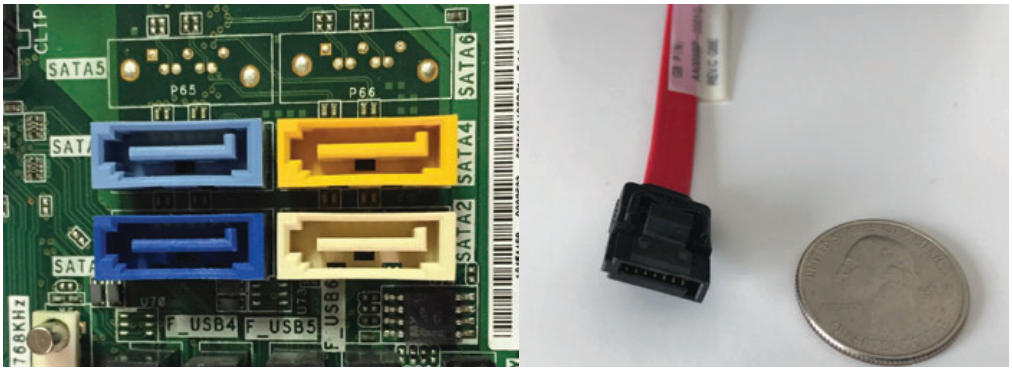

The SATA we've discussed so far is internal, but there's an external version as well, appropriately named *external SATA (eSATA)*. It uses the same technology, only in an external connection. The port at the bottom center of the left image in Figure 3.31 is eSATA, and the image on the right is an eSATA cable. It entered the market in 2003, is mostly intended for hard drive use, and can support up to 15 devices on a single bus.

Table 3.4 shows some of the eSATA specifications.

TABLE 3.4 eSATA specifications

Version	Year	Speed	Names
Revision 1.0	2003	1.5 Gbps	SATA I, SATA 1.5 Gb/s
Revision 2.0	2005	3.0 Gbps	SATA II, SATA 3 Gb/s
Revision 3.0	2009	6.0 Gbps	SATA III, SATA 6 Gb/s

FIGURE 3.31 eSATA port and cable

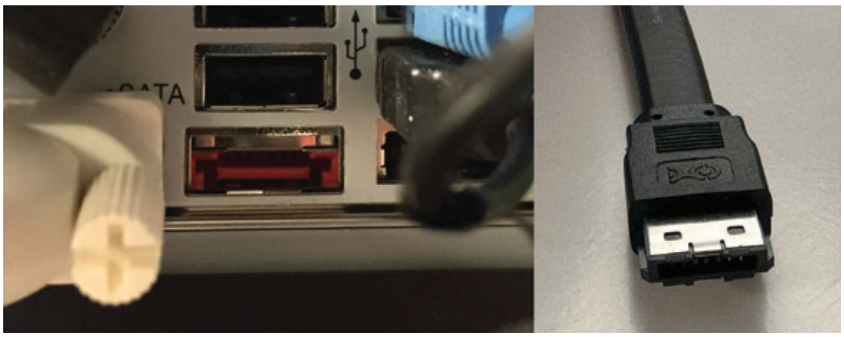

You will commonly see the third generation of eSATA (and SATA) referred to as SATA 6 or SATA 6 Gb/s. This is because if they called it SATA 3, there would be confusion with the second generation, which had transfer speeds of 3.0 Gbps.

An interesting fact about eSATA is that the interface does not provide power, which is a big negative compared to its contemporary high-speed serial counterparts. To overcome this limitation, there is another eSATA port that you might see, called *Power over eSATA*, eSATA+, eSATAp, or eSATA/USB. It's essentially a combination eSATA and USB port. Since the port is a combination of two others, neither sanctioning body officially recognizes it (which is probably why there are so many names—other companies call it what they want to). Figure 3.32 shows this port.

FIGURE 3.32 USB over eSATA

You might have noticed that the most recent eSATA specification was released over a decade ago. eSATA is basically dead, because newer standards such as USB4 and Thunderbolt 3 (and newer) provide not only power but faster data transmission speeds as well.

You can see that this port is slightly different from the one in Figure 3.31, and it's also marked with a USB icon next to the eSATA one. On the market, you can purchase cables that go from this port to an eSATA device and provide it with power via the eSATAp port.

A Curtain Call for PATA

We mentioned PATA in Volume 1 Chapter 2 and noted that it's essentially obsolete now, but you still might run into it on occasion. In addition, the power connector for PATA drives (and older optical drives) is the *Molex connector*, which still does happen to be an exam objective. So for the sake of completing our review of hard drive cables and connector types covered in the exam objectives, Figure 3.33 shows a 4-pin Molex connector. Historically, most were white or opaque plastic, but there's no standard color.

FIGURE 3.33 Molex connector

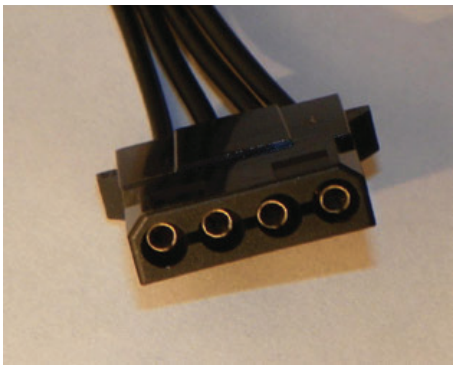

Summary

In this chapter, you first learned about peripheral types. We broke them into four categories: video, audio, input/output, and storage. Video peripherals include monitors and projectors. There aren't many audio connectors, but most use the TRS connector you learned about. Input and output devices are plentiful, and we concentrated on keyboards and mice. Storage devices and optical drives can be external, and an example is an external network-attached storage (NAS) device.

In the second section of the chapter, you learned about various cable and connection types, and the purposes and uses of peripheral types. First, you learned about peripheral cables and connectors, such as USB and its variants, Lightning, Thunderbolt, and serial. Then we moved on to video cables. Topics included the analog VGA standard, as well as the digital standards DVI, HDMI, and DisplayPort. Then, we covered hard drive connections and cables related to SATA and eSATA, with a quick nod to the PATA-era Molex power connector.

Exam Essentials

Understand the differences between display types. You should be able to know the advantages and disadvantages of using IPS, TN, and VA LCD displays, as well as OLEDs and mini-LEDs.

Know what touch screens, digitizers, and inverters are used for. Touchscreens and digitizers have the same function, which is to take physical input from a finger or stylus and convert that to digital signals for the computing device to use. An inverter is a small circuit board that converts DC power to AC power for an LCD backlight.

Know four different display attributes. Be familiar with pixel density (the number of pixels per inch), refresh rates (how often the screen is redrawn), screen resolution (the number of pixels across and down on a display), and color gamut (the range of colors a display can produce).

Recognize and understand different peripheral connectors and adapters. Expansion cards and motherboards have external connectivity interfaces. The interfaces have connectors that adhere to some sort of standard for interconnecting with a cable or external device. Knowing these specific characteristics can help you differentiate among the capabilities of the interfaces available to you. Understanding when to use an adapter to convert one connector to another is crucial to achieving connectivity among differing interfaces. Adapters you should know are DVI to HDMI, USB to Ethernet, and DVI to VGA.

Recognize and be able to describe display connectors specifically. Although a type of peripheral connector, display connectors are in a class all their own. Types include VGA, various versions of DVI, HDMI, mini-HDMI, DisplayPort, and USB-C.

Recognize and understand the purpose of hard drive cables and connectors. The connectors you should recognize are SATA and eSATA. Molex connectors are used to power older PATA devices. They each connect hard drives or optical drives.

Know the various peripheral cables and their connectors. Multipurpose cables include USB, Lightning, Thunderbolt, and serial. USB has the largest variety of connectors, including USB-A and USB-B and their mini- and micro- versions, as well as the newer USB-C. USB cables also can have proprietary connectors, such as Apple's Lightning connector. Thunderbolt can use a proprietary connector or USB-C. Serial cables have a DB-9 connector.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. Which of the following could be used with an analog monitor? (Choose two.)
   - **A.** HDMI
   - **B.** DVI-A
   - **C.** DVI-D
   - **D.** DisplayPort
   - **E.** VGA

2. You are attaching an external hard drive to a computer with USB 2.0 ports. The user is concerned about data transfer speed. What is the maximum speed of USB 2.0?
   - **A.** 480 Mbps
   - **B.** 5 Gbps
   - **C.** 10 Gbps
   - **D.** 20 Gbps

3. You got a new desktop computer for a designer at your company. On the front panel is a rectangular yellow USB port. What does that mean?
   - **A.** It's USB 2.0.
   - **B.** It's USB 3.0.
   - **C.** It's a port capable of charging devices even with the system powered off.
   - **D.** It's a port to be used with USB-C compatible devices.

4. You are installing a replacement monitor on a desktop computer. Which of the following is true regarding a monitor's refresh rate?
   - **A.** As long as the graphics adapter can refresh the image at a particular rate, the attached monitor can accommodate that refresh rate.
   - **B.** The refresh rate is normally expressed in MHz.
   - **C.** The refresh rate is normally selected by using the controls on the front panel of the monitor.
   - **D.** As you lower the resolution, the maximum refresh rate allowed tends to increase.

5. Which of the following connector types is typically not associated with hard drive use?
   - **A.** USB-C
   - **B.** eSATA
   - **C.** DB-9
   - **D.** SATA

6. The company photographer is interested in buying a new digital camera. He wants to be able to plug it directly into his laptop to transfer pictures for editing. What type of connector will most likely be on his camera?
   - **A.** Mini-USB
   - **B.** Thunderbolt
   - **C.** Molex
   - **D.** Lightning

7. You are installing a new video card that has an HDMI port. Which of the following are true about HDMI cables? (Choose all that apply.)
   - **A.** They can transmit analog or digital signals.
   - **B.** They can transmit audio as well as video.
   - **C.** They are not used to connect to projectors.
   - **D.** They can use HDMI or mini-HDMI connectors.

8. Which of the following display interfaces is equivalent to DisplayPort with PCIe added in?
   - **A.** Thunderbolt
   - **B.** HDMI
   - **C.** DVI
   - **D.** VGA

9. Which of the two standards use the same connector? (Choose two.)
   - **A.** eSATA
   - **B.** Thunderbolt 3
   - **C.** USB4
   - **D.** USB 2.0
   - **E.** Mini-USB

10. How do you connect a DVI-A interface on a peripheral to a DVI-D interface on the computer?
    - **A.** With a DVI-I cable.
    - **B.** With a cable that is terminated on one end with a DVI-A connector and on the other end with a DVD-D connector.
    - **C.** You wouldn't interconnect those two interfaces.
    - **D.** With a standard DVI cable.

11. A technician says he needs to provide power to a PATA hard drive and a DVD-ROM in a desktop. The computer now has a new fully modular power supply, and the technician wants to know what type of power connectors he should take. What do you tell him?
    - **A.** SATA
    - **B.** SCSI
    - **C.** Serial
    - **D.** Molex

12. Which of the following connectors is oval shaped?

A. Micro-USB

B. Mini-USB

C. Thunderbolt 2

D. Thunderbolt 3

13. Which of the following connector types are capable of conducting power and transmitting data? (Choose all that apply.)

A. Lightning

B. Molex

C. USB-C

D. DB-9

E. Micro-USB

14. Which of the following statements is true about eSATA cables and connectors?

A. eSATA cables provide power and data to devices.

B. eSATA connectors are smaller than USB-C connectors.

C. eSATA is faster than USB4 and Thunderbolt 3.

D. eSATA maximum throughput is 6 Gbps.

15. Which of the following is not an example of a connector for a standard peripheral input device?

A. 1/8″ jack

B. Molex

C. D-sub

D. USB

16. You are asked to hook up a projector in a classroom. What type of cable are you most likely to need?

A. USB Type A

B. eSATA

C. HDMI

D. DB-9

17. A user with an iPhone 14 needs to charge her device and asks you to bring a cable. Which of the following connector types should you ensure that the cable has?

A. Thunderbolt

B. Lightning

C. Mini-USB

D. Micro-USB

18. Which of the following types of cables will transfer data at the slowest rate?

A. USB4

B. Serial

C. Thunderbolt 3

D. HDMI

19. What is the maximum data transfer rate of USB High Speed?

A. 10 Gbps

B. 5 Gbps

C. 480 Mbps

D. 12 Mbps

20. VGA cables use what type of signal between the adapter and monitor?

A. Digital

B. Analog

C. Composite

D. Compressed

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answers compare to the authors', refer to Appendix B.

Looking at the back of a computer, you see the interfaces shown in the following graphic. Which type of cables do you need to plug into each one?

A:

B:

C:

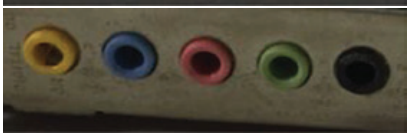

D:

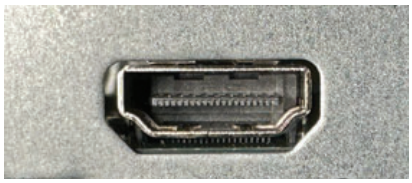

E:

F:

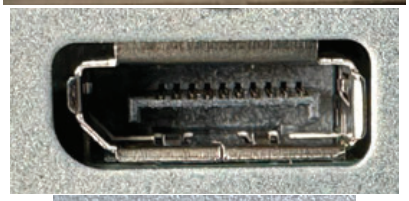

G:

H:

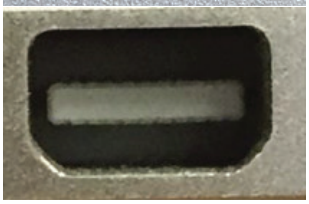

Chapter

4

Laptop and Mobile Device Hardware

THE FOLLOWING COMPTIA A+ 220-1201 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

✓ **1.0 Mobile Devices**

✓ **1.1 Given a scenario, monitor mobile device hardware and use appropriate replacement techniques.**

- Battery
- Keyboard/keys
- Random-access memory (RAM)
- Hard disk drive (HDD)/solid-state drive (SSD)
- Wireless cards
- Physical privacy and security components
  - Biometrics
  - Near-field scanner features
- Wi-Fi antenna connector/placement
- Camera/webcam
- Microphone

✓ **1.2 Compare and contrast accessories and connectivity options for mobile devices.**

- Connection methods
  - Universal Serial Bus (USB)/USB-C/microUSB/miniUSB
  - Lightning
  - Near-field communication (NFC)
  - Bluetooth
  - Tethering/hotspot

- Accessories
  - Stylus
  - Headsets
  - Speakers
  - Webcam
- Docking station
- Port replicator
- Trackpad/drawing pad/track points

In 1943, the president of IBM, Thomas Watson, was quoted as saying, "I think there is a world market for maybe five computers." Somewhat more recently, in 1977, Ken Olsen, the founder of one-time computer industry giant Digital Equipment Corporation (DEC), stated, "There is no reason anyone would want a computer in their home." Four years later the personal computer was introduced, ironically by IBM. It's about 80 years past the first quote and about 50 years past the second, but looking back at the history of computers, it's amazing to see how far the industry has come.

As recently as the 1990s, portable computers were luxuries that were affordable to only the wealthy or the select few businesspeople who traveled extensively. As with all other technologies, though, portable systems became smaller, lighter (more portable), more powerful, and less expensive. Because the technology and price disparity between the two platforms has decreased significantly, laptops have outsold desktops since the mid-2000s. For about 20 years, laptop computers enjoyed a status as the most popular mobile computing device available. Desktop computers still were more powerful and cheaper, but with a laptop, users weren't tethered to a desk.

Technology continued its inevitable progression, and smaller mobile devices came into play. Tablets, smartphones, and even wearable devices obtained enough power and features to be considered computers. Laptop sales have been flat since about 2012, whereas smartphones and wearables have continued to increase in market share. Further, smartphones overtook laptops as the most popular device to access the Internet as of 2015. In stark contrast to the computers of yesteryear, today's most popular computing devices can be so small that the likelihood of misplacing them is a real and ongoing concern.

Every indication is that the movement toward mobile computing will continue, so you definitely need to be well versed in portable technologies, which contain both nifty features and frustrating quirks. In this chapter, you will learn about the features and quirks of laptop and mobile device hardware. There are some differences—for example, the majority of hardware in a laptop is more similar to that of a desktop than to that of a smartphone, but there are similarities as well. The display of a laptop (and to be fair, most desktop-sized monitors) is a lot like those in smaller devices. Throughout the chapter, we will cover laptops and smaller mobile devices as though they are similar, and call out specific differences where appropriate. We will start with installing and configuring laptop hardware and components, including details on display technologies. Then we'll finish with setting up and configuring accessories and ports on mobile devices.

Working with Laptop and Mobile Device Hardware

Hardware in all computing devices needs to perform the same tasks. There are devices that control input and output, processing, short-term and long-term storage, displaying information, and connecting to other computers. This is true regardless of the size of the device. Granted, smaller devices are space-constrained, so the hardware components that do these tasks need to be smaller and consume far less energy.

Working with hardware, then, should be similar across devices, and for the most part it is. Maybe instead of having a physical keyboard you'll have a virtual one (and you can't exactly replace it, per se), but at least you know how to input to the device. Or maybe the display is smaller, but it and the video card still perform the same tasks that their larger counterparts do. In the following sections, we will look at the specifics regarding working with hardware in laptops and smaller devices.

Understanding the Differences Between Device Types

The first personal computers developed were more similar to today's desktop computers than they were to laptops or mobile devices. The smaller devices do trace their ancestry back to the desktop, though, so it makes sense to compare them to their bigger counterparts. Here, we'll take a high-level look at what makes laptops and mobile devices unique.

Laptops vs. Desktops

Laptops are similar to desktop computers in architecture in that they contain many parts that perform similar functions. However, the parts that make up a laptop are completely different from those in desktop computers. The obvious major difference is size; laptops are space challenged. Another primary concern is heat. Restricted space means less airflow, meaning parts can heat up and overheat faster.

To overcome space limitations, laptop parts are physically much smaller and lighter, and they must fit into the compact space of a laptop's case. It might not sound like much, but there really is a major difference between a 4-pound laptop and a 5-pound laptop if you're hauling it around in its carrying case all day. Also, laptop parts are designed to consume less power and to shut themselves off when not being used, although many desktops also have components such as video circuitry that go into a low-power state when not active. Finally, most laptop components are proprietary—not just to the manufacturer but even to the specific model. The motherboard is especially proprietary, and the *liquid crystal display (LCD)* screen from one laptop will not necessarily fit on another.

Manufacturers have also pushed out smaller and smaller portables that are laptop adjacent. For example, in 2007 the first netbooks were introduced. A *netbook* is an extremely small laptop computer that is lighter in weight and more scaled down in features than a standard laptop. The term *netbook* is rarely used today, but Chromebooks are an example of that type of technology. Users are attracted to Chromebooks because of their enhanced portability and affordability. The features that remain are ideal for Internet access and emailing.

However, many users would find them insufficient for mainstream business usage. Tablets are even smaller yet, but they are typically characterized as a mobile device, so we'll hold off on talking about them just yet.

If you've shopped for a laptop, you have no doubt noticed that the prices of desktop PCs are often quite a bit lower than those for laptop computers, yet the desktops are usually faster and more powerful. If you've ever wondered what makes a laptop so much different from a PC, here are the primary differences between laptops and desktops:

Portability This is probably the most obvious difference. Laptops are designed to be portable. They run on batteries, so you aren't tied to one spot at home or at the office. Networking options are available that allow you to connect to a network wirelessly and do work from just about anywhere, including airports, coffee shops, and so on. As anyone who has tried to bring their full-tower PC to a LAN party can tell you, desktops just aren't that portable.

Cost Laptops tend to cost more than desktop computers with similar features. The primary reason is that portability requires smallest components and unique proprietary designs so that those components fit into the smallest size necessary. Miniature versions of components cost more money than standard-sized (desktop) versions. The cost discrepancy between desktops and laptops has shrunk considerably in the last several years, but it still exists.

Performance By and large, laptops are always going to lose out somewhere in the performance department. Compromises must often be made between performance and portability, and considering that portability is the major feature of a laptop, performance is what usually suffers. While it is possible to have a laptop and a desktop with comparable performance, the amount of money one would have to spend for a "desktop replacement" laptop is considerable. This is not to say that a laptop can't outperform a desktop—it's just that the "bang for the buck" factor is higher in a desktop.

Expandability Because desktop computers were designed to be modular, their capabilities can be upgraded quite easily. It is next to impossible to upgrade the processor or motherboard on most laptops. Other than memory and hard drives, most laptop upgrades consist of adding an external device through one of the laptop's ports, such as a USB port.

Quality of Construction Considering how much abuse laptops get, it is much more important that the materials used to construct the laptop case and other components be extremely durable. Durability is important in a desktop too, but it won't be tested as much as in a laptop.

Building Your Own

This anecdote comes from one of the authors: "During an A+ course, I gave the class the assignment to go out on the web and put together the most powerful and complete computer they could for under a thousand dollars. The class was for non-degree-seeking adults,

so nothing was graded; it was simply to provide experience with spec'ing out and pricing the parts that go into making a complete system.

"One of the students had her eye on a new laptop for personal use. Because she noticed the trend toward being able to build a desktop computer for less than she could buy one, the student assumed the same about laptops. Unfortunately, I had not specifically mentioned the fact that there are no standards for building complete laptop clones, unlike with desktops.

"You can't reliably build your own laptop. Because laptop components are designed to exacting specifications to fit properly inside one manufacturer's case, there generally are no universal motherboards, video boards, and so on for laptops. Memory and hard drives are the exception. You can get different brands of memory and hard drives for laptops, but you can't buy a motherboard from one company and the video circuitry from another."

Characteristics of Mobile Devices

If you were asked to define the primary characteristic of mobile devices, you would probably answer, "They are small," and you wouldn't be wrong. Three overarching characteristics of mobile devices make working with them unique versus working with laptops or desktops: field servicing and upgrading, input methods, and secondary storage. We'll discuss each one in turn.

Field Servicing and Upgrading

Ever since the dawn of the portable computer, manufacturers and service providers have based a percentage of their success on warranties and "house calls" to repair devices on the fritz. It's a fact that quasi-permanent components, such as displays and motherboards, are widely considered replaceable only with identical components in laptops and smaller devices. However, technically minded users could take it upon themselves to expand the capabilities of their own system by, for instance, upgrading the hard drive, increasing RAM, using expansion cards and flash devices, and attaching wired peripherals.

Although the ability to repair and expand the functionality of portable devices in the field has become all but obviated, it has been shown with current and past generations of mobile devices that users are not averse to giving up expandability and replaceable parts as long as functionality and convenience outshine the loss.

Although many Android and other non-Apple devices allow the replacement of batteries and the use of removable memory cards as primary storage, even this basic level of access is removed in Apple's mobile devices, including its iPad line of tablet computers. In an effort to produce a sleeker mobile phone, even Android devices have been developed without user access to the battery. For Apple, however, in addition to producing a nice compact package, it is all part of keeping the technology as closed to adulteration as possible. Supporters of this practice recognize the resulting long-term quality. Detractors lament the lack of options.

To service closed mobile devices of any size, you may have to seek out an authorized repair facility and take or send your device to them for service. Attempting your own

repairs can void any remaining warranty, and it can possibly render the device unusable (also known as "bricking" a phone, because you basically limited its functionality to that of a brick). For example, a special screwdriver-like tool is required to open Apple's devices. You cannot simply dig between the seams of the case to pop the device open. Even if you get such a device to open, there is no standard consumer pipeline for parts, whether for repair or upgrading. If you want to try the repair yourself, you could be on your own. You may be able to find helpful videos on YouTube or `www.ifixit.com` to provide some guidance, though.

Anyone who has been around the business for more than just a few years has likely seen their fair share of components and systems with no user-serviceable parts. For these situations, an authorized technician can be dispatched to your location, home or work, with the appropriate tools, parts, and skills to field-service the system for you. On a slightly different, perhaps subtler note, the bottom line here is that many of today's mobile devices, including some of the larger tablet-style devices, have no field-serviceable parts inside, let alone user-serviceable parts. In some extremes, special work environments similar to the original clean manufacturing environment have to be established for servicing.

Input Methods

With decreased size comes increased interaction difficulties. Human interfaces can become only so small without the use of projection or virtualization. In other words, a computer the size of a postage stamp is fine as long as it can project a full-sized keyboard and a 60″ display, for example. Using microscopic real interfaces would not sell much product. Thus, the conundrum is that users want smaller devices, but they do not want to have to wear a jeweler's loupe or big-screen virtualization glasses to interact with their petite devices.

As long as the size of the devices remains within the realm of human visibility and interaction, modern technology allows for some pretty convenient methods of user input. Nearly all devices, from tablet size down, are equipped with touchscreens, supplying onscreen keyboards and other virtual input interfaces. On top of that, more and more of the screens are developing the capability to detect more than one contact point.

We discuss touchscreens more in the section "Input Devices," later in this chapter.

Generically, this technology is referred to in the industry as multitouch, and it is available on all Apple devices with touch input, including the touchpads of the Apple laptops. Apple, through its acquisition of a company called FingerWorks, holds patents for the capacitive multitouch technology featured on its products. Today, multitouch is more about functionality than novelty. Nevertheless, the markets for both business and pleasure exist for multitouch.

Certainly, touchscreens with the capability to sense hundreds of separate points of contact can allow large-scale collaboration or fun at parties. Imagine a coffee table that can allow you to pull out a jigsaw puzzle with the touch of an icon, remembering where you and three friends left off. Imagine all of you being able to manipulate pieces independently and simultaneously and being able to send the puzzle away again as quickly as you brought it out so that you can watch the game on the same surface. This technology exists, and it is for

sale today. Early examples were built on Microsoft's PixelSense technology, including the Samsung SUR40. Companies like Ideum build multitouch platform tables, including a monster 86″ 4K ultra-high-definition (UHD) display with 100 touchpoints, allowing up to eight people to use it simultaneously.

On a smaller scale, our mobile devices allow us to pinch and stretch images on the screen by placing multiple fingers on that screen at the same time. Even touchpads on laptops can be made to differentiate any number of fingers being used at the same time, each producing a different result, including pointing and clicking, scrolling and right-clicking, and dragging—all one-handed with no need to press a key or mouse button while gesturing.

HTC created an early touchscreen software interface called TouchFLO that has matured into HTC Sense, and it is still in use today on its Android-based mobile devices. TouchFLO is not multitouch capable, nor does it specify the physical technology behind the touchscreen, only the software application for it. Theoretically, then, TouchFLO and multitouch could be combined.

The primary contribution of TouchFLO was the introduction of an interface that the user perceives as multiple screens, each of which is accessible by an intuitive finger gesture on the screen to spin around to a subsequent page. On various devices using this concept, neighboring pages have been constructed side by side or above and below one another. Apple's mobile devices employ gestures owing to the contributions of TouchFLO, bringing the potential of combining TouchFLO-like technology and multitouch to bear.

Users of early HTC devices with resistive touchscreen technology met with difficulty and discord when flowing to another screen. The matte texture of the early resistive screens was not conducive to smooth gesturing. Capacitive touchscreen technology is a welcome addition to such a user interface, making gestures smooth and even more intuitive than ever.

Recall from Volume 1 Chapter 3, "Peripherals, Cables, and Connectors" that resistive touchscreens use pressure and capacitive touchscreens use a change in electrical conductivity to detect a touch.

Secondary Storage

Computers of all sizes and capabilities use similar forms of RAM for primary storage—the storage location for currently running instructions and data. Secondary storage—the usually nonvolatile location where these instructions and data are stored on a more permanent basis—is another story.

The primary concern with smaller devices is the shock they tend to take as the user makes their way through a typical day. Simply strapping a phone to your hip and taking the metro to work presents a multitude of opportunities for a spinning disk to meet with catastrophe. The result would be the frequent loss of user information from a device counted on more and more as technology advances.

Just as most telephony subscribers have migrated from a home landline that stays put to a mobile phone that follows them everywhere, many casual consumers are content to use their mobile device as their primary or only computing system, taking it wherever they go. As a result, the data must survive conditions more brutal than most laptops because laptops are most often shut down before being transported.

The most popular solution is to equip mobile devices with very small solid-state drives (SSDs) in place of larger magnetic or solid-state drives. There are no moving parts, the drive stays cooler and resists higher temperature extremes, and SSDs require less power to run than their conventional counterparts.

Now that we've illustrated the primary differences between laptops, mobile devices, and desktops, let's examine some principles for taking laptops apart and putting them back together.

Disassembling and Reassembling Laptops

Desktop computers often have a lot of empty space inside their cases. This lets air circulate and also gives the technician some room to maneuver when troubleshooting internal hardware. Space is at a premium in laptops, and rarely is any wasted. With a desktop computer, if you end up having an extra screw left over after putting it back together, it might not be a big deal. With laptops, every screw matters, and you'll sometimes find yourself trying to identify visually miniscule differences between screws to make sure that you get them back into the right places.

Even though repairing a laptop poses unique issues, most of the general troubleshooting and safety tips that you use when troubleshooting a desktop still apply. For example, always make sure that you have a clean and well-lit workspace and be cautious of electrostatic discharge (ESD). General safety tips and ESD prevention are covered in Volume 2 Chapter 9, "Safety and Environmental Concerns." For now, our general advice is to use antistatic mats or wrist straps if they're available.

Throughout this section, we'll use the word *laptop* almost exclusively. The principles covered here apply to nearly all portable devices, though, even tablets and smartphones.

One of the key principles for working with laptops is using the right tools to tear the thing apart. It's doubtful that any technician goes into a job thinking, "Hey, I'm going to use the wrong tools just to see what happens." With laptops, though, it's especially important to ensure that you have exactly the tools you'll need for the job. The two main camps of materials you need are the manufacturer's documentation and the correct hand tools. We'll also emphasize the importance of organization.

Using the Manufacturer's Documentation

Most technicians won't bat an eye at whipping out their cordless screwdriver and getting into a desktop's case. The biggest difference between most desktops is how you get inside the case. Once it's opened, everything inside is pretty standard fare.

Laptops are a different story. Even experienced technicians will tell you to not remove a single screw until you have the documentation handy and have taken pictures of the pre-surgery state (if you're allowed to) unless you're incredibly familiar with that particular

laptop. Most laptop manufacturers give you access to repair manuals on their websites. Table 4.1 lists the service and support websites for some of the top laptop manufacturers.

TABLE 4.1 Laptop manufacturers' service and support websites

Company	URL
Apple	`https://support.apple.com/mac`
Asus	`https://www.asus.com/support`
Dell	`https://www.dell.com/support`
HP	`https://support.hp.com`
Lenovo	`https://support.lenovo.com`
Sony	`https://www.sony.com/electronics/support`

If the site you need isn't listed, a quick Google search should do the trick. Once you are at the right website, search for the manual using the laptop's model number.

Some laptop manufacturers have a policy that if you open the case of a laptop, the warranty is voided. Be sure to understand your warranty status and implications of cracking the case open before you do it.

Using the Proper Hand Tools

Once you have the manual in hand or on your screen, you must gather the proper hand tools for the job. For some laptops, you only need the basics, such as small Phillips-head and flat-head screwdrivers. For others, you may need a Torx driver. Gather the tools you need and prepare to open the case. A small flashlight might also come in handy. Small PC technician toolkits are readily available online or from your favorite electronics retailer. They may have a few different sizes of screwdrivers, hex drivers, Torx drivers, tweezers, a screw grabber, and a few other assorted goodies, all in a convenient carrying case. An example is shown in Figure 4.1. Find one you like and never leave home without it.

The Consequences of Using the Wrong Tools

It's been said once, but it's important to say it again: always use the right tool for the job when repairing laptops. If the documentation says that you need a T-10 Torx driver, make sure that you have a T-10 Torx driver.

Not using the right tools can result in the stripping of the screw head. If you strip a screw head in a desktop, you might have alternative methods of removing the screw. Laptops are far less forgiving. If you strip a screw head and are unable to turn the screw, you may never be able to remove it. That could result in needing to scrap the device.

FIGURE 4.1 PC technician toolkit

Organization and Documentation

Before you crack open the case of your laptop, have an organization and documentation plan in place. Know where you are going to put the parts. Have a container set aside for the screws. You can purchase small plastic containers that have several compartments in them with lids that snap tightly shut, into which you can place screws. You can also use containers designed to organize prescription pills or fishing tackle. The bottom of an egg carton works well too, provided that you don't need to transport the screws from place to place. You don't want the screws falling out and getting lost!

For documentation, many technicians find it handy to draw a map of the computer they're getting into, such as the one shown in Figure 4.2. It can be as complex as you want it to be, as long as it makes sense to you. Taking pictures with your phone is also a smart move, provided that you're allowed to use your phone and don't violate any security or privacy policies.

The drawing in Figure 4.2 shows the locations of the screws, and it also calls out where the screws should be placed once they're removed. Again, this type of documentation can be as simple or complex as you want it to be, as long as it makes sense and helps you stay organized.

Now that we've covered some key principles, let's take a look at specific components, technologies involved, and how to install and configure them.

FIGURE 4.2 Laptop repair road map

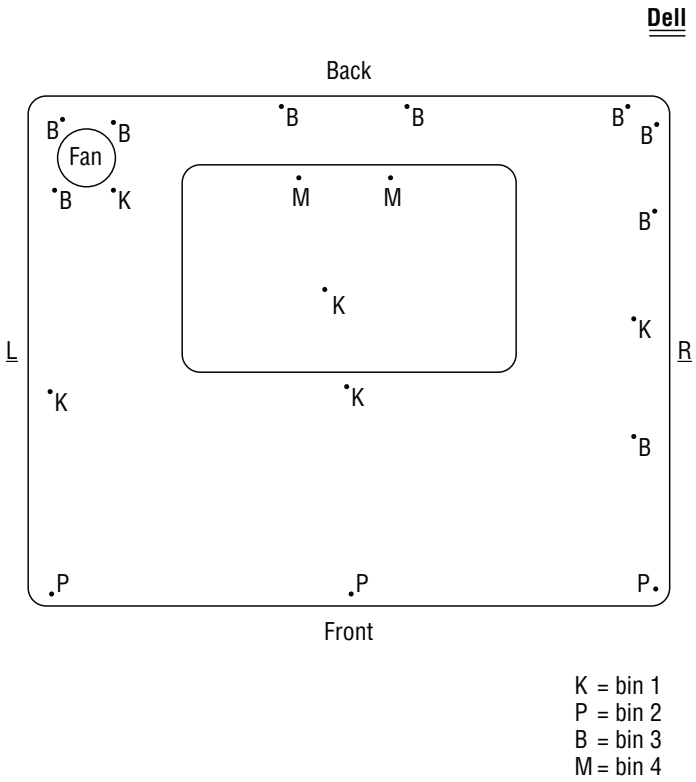

Installing and Configuring Laptop Hardware

In the following sections, you will learn about the various components that make up laptops and how they differ from desktop computer components. These sections deal specifically with laptops, because smaller devices generally don't have field-replaceable components (or you can get specialized training on how to repair them). If you don't remember exactly what each component does, it may help you to refer back to earlier hardware chapters occasionally as you read this chapter.

Laptop Case

A typical laptop case is made up of three main parts:

- The display—usually an LCD, light-emitting diode (LED), or organic light-emitting diode (OLED) display
- The case structure, which is the metal reinforcement inside the laptop that provides rigidity and strength and to which most components are mounted
- The case, or the plastic or aluminum cover that surrounds the components and provides protection from the elements

Some laptops have cases made of a strong, lightweight metal, such as aluminum or titanium. However, the majority of laptop cases are made of plastic, usually acrylonitrile butadiene styrene (ABS) or an ABS composite.

Laptop cases are made in what is known as a *clamshell design*. In a clamshell design, the laptop has two halves, hinged together at the back. The display portion, called the top half, often includes a webcam, microphone, and Wi-Fi antenna. All other components, including the motherboard, memory, storage, keyboard, battery, cooling fan, and speakers, are in the bottom half.

Occasionally, part of the laptop's case or the device's frame will crack and need to be replaced. However, you usually can't just replace the cracked section. Most often, you must remove every component from inside the laptop's case and swap the components over to the new one. This is a labor-intensive process because the screws in laptops are often very small and hard to reach.

Often, repairing a cracked case may cost several hundred dollars in labor alone. Most times, people who have cracked laptop cases wait until something else needs to be repaired before having the case fixed. Or, they just wait until it's time to upgrade to a new system. The decision on when to repair or replace the laptop boils down to a few factors. The primary one is if the user can live with the damage. While they can be annoying, most case problems don't inhibit the operation of the machine. The secondary factor is money. The user (or company) needs to decide if it's really worth spending the money needed to fix the issue immediately.

One of the components you may need to service is the speakers, which are generally built into the bottom part of the case. Exercise 4.1 will provide an example of removing speakers from a Dell Inspiron 13 7000 laptop. We'll use that model for an example throughout this chapter. We realize that this specific model is a few years old and you might not have access to one, but the exercises can help you understand the process. Besides, there are two key points regarding these exercises:

- The basic clamshell laptop design hasn't changed in over three decades. Parts have gotten smaller but they are likely in similar places to where their predecessors were.
- It would be impractical to list the specific replacement steps for all makes and models out there. If you have a different model, download its service manual and perform the necessary steps.

To get the most out of the exercises in this chapter, it's best if you're able to perform them on a laptop. However, we realize that not everyone has access to a laptop they can tear apart. If you don't have a laptop you can perform surgery on, here are a few options you can try:

- Download a service manual for the laptop you have (or another model that seems interesting). Look for directions on how to remove or replace the components we describe so that you can get a feel for how you would perform the tasks.
- Visit a local computer repair shop and ask if they have any old laptops that they are going to take to recycling. Explain that you are studying for the A+ exam and would like to get experience taking apart a laptop. See if they will let you take one apart at their facility. (You never know until you ask!)
- Visit YouTube or `www.ifixit.com` and look up videos for how to replace components on your laptop or another model that seems interesting.

Before we get to Exercise 4.1, we want to remind you of a few safety steps to take before working on a laptop. We're not going to repeat these instructions before every exercise, but always perform these steps before beginning digging into a laptop (or a desktop, for that matter):

1. Turn off the computer.
2. Disconnect all external peripherals and cables, such as keyboards, mice, monitors, and network cables.
3. Unplug it from the power source.
4. Ensure that your workspace is free of clutter and well lit.
5. Have the manual or help videos handy, just in case.
6. Make sure you have the proper tools readily available. (This might include a phone to take pictures with.)
7. Ground yourself with an ESD wrist strap or other antistatic protection.

EXERCISE 4.1

Removing Speakers from a Laptop

1. Remove the bottom cover. As shown in Figure 4.3, this model requires removing nine screws to take the cover off.

FIGURE 4.3 Removing nine screws to take off the bottom cover

2. Make a mental note of the key components and their location, or better yet, take a picture with your phone like the one in Figure 4.4.

FIGURE 4.4 Inside the laptop

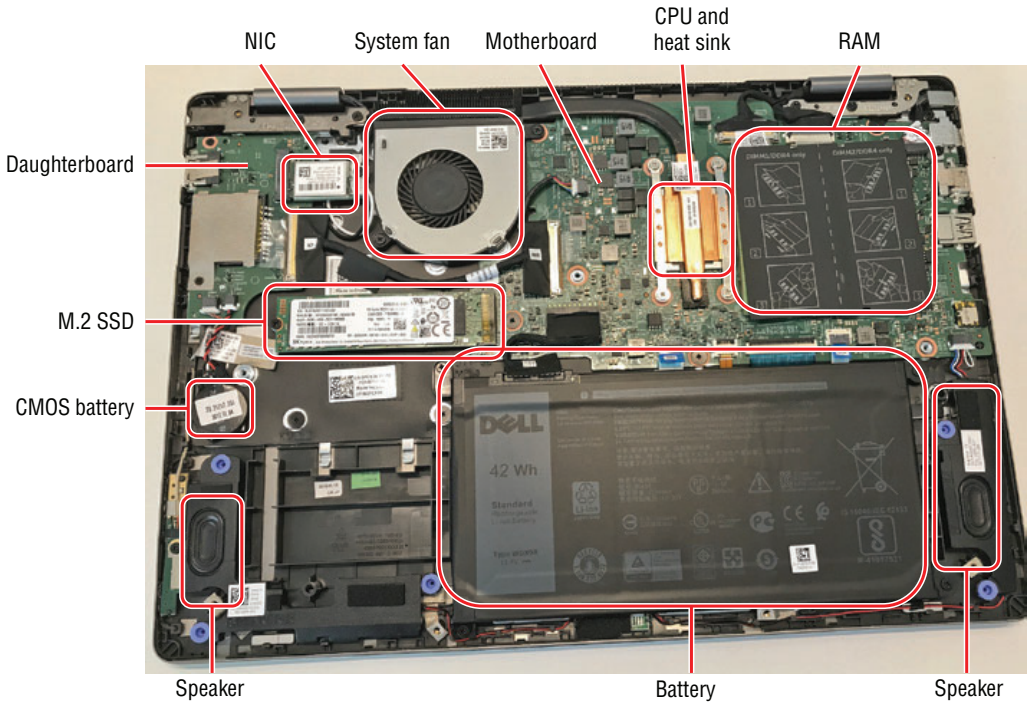

3. Disconnect the speaker cable from the motherboard, as shown in Figure 4.5.

FIGURE 4.5 Disconnecting the speaker cable

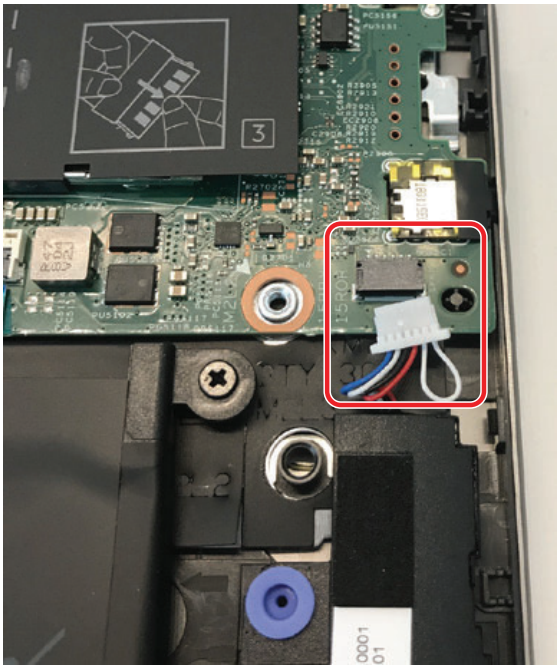

EXERCISE 4.1 *(continued)*

4. Remove the speaker cable from the wire routing guides on the base assembly.

   Figure 4.6 shows the speaker wire that travels from the left speaker to the right speaker.

FIGURE 4.6 The speaker wire

5. Remove the rubber grommets that hold the speakers in.

   Figure 4.6 shows the two from the right speaker. The left speaker has three.

6. Carefully remove both speakers from the base assembly.

To replace the speakers, follow the steps in reverse order. Insert the new speakers, replace the grommets, run the speaker wire through the guides, and connect the speaker to the motherboard.

Removing the speakers from the Dell Inspiron 13 7000 was relatively easy. On some laptops, the speakers or speaker wire are completely buried underneath several other components, and you need to remove those first. As we move through this chapter, you will see situations where you need to remove several components before you can get to the one you want.

Laptop Displays

The display system is the primary component in the top half of the clamshell case. (The wireless antenna, webcam, and microphone often reside here too, and we'll get to them in just a bit.) Much like all other laptop components, the display is more or less a smaller version of its desktop counterpart. What is unique to laptop displays, though, is that for some

time, the technology used in them was actually more advanced than what was commonly used in desktops. This is due to LCD technology.

Before LCD technology, computer displays used cathode-ray tube (CRT) technology (like old-school televisions) and were big, bulky, and hardly mobile. We introduced LCD, LED, and OLED concepts in Volume 1 Chapter 3, but we'll do a quick refresher in the "Screen" section after we talk about video cards. Let's quickly cover the different components that are required to make these types of displays work on a laptop.

Some laptops feature screens that rotate orientation depending on how they are held (much like smartphones and tablets), or the screens themselves will rotate horizontally or be entirely removable. Rotating and removable screens cost more money, so consider your needs or the needs of the user when recommending such features.

Video Card

The video card in a laptop or desktop with a flat-screen (LCD, LED, OLED, or Mini-LED) monitor does the same thing, regardless of what type of machine it's in. It's responsible for generating and managing the image sent to the screen. In this section we will bundle all of those types of monitors together as LCDs for the sake of brevity. LCD monitors are digital, so laptop video cards generate a digital image. Laptop manufacturers put video cards that are compatible with the given display in a laptop, and most laptop manufacturers choose to integrate the LCD circuitry on the motherboard to save space.

Odds are that the laptop you're using has an integrated video card. If the video card fails, you're likely looking at a motherboard replacement. Very few laptops have a replaceable video card. For those that do, replacing or upgrading it will probably resemble replacing system memory.

Real World Scenario

Video Memory Sharing

If your video card is built into your motherboard, it likely won't have its own memory but will share system memory with the processor. Note that there is nothing wrong with this type of setup; in fact, it often brings the cost of the laptop down. It's just that instead of having 16 GB of RAM and 4 GB of video RAM (for example), you would have only 16 GB total. So if your video card were using 4 GB, the system would be left with only 12 GB. That will slow everything down.

How much of a difference does all of this make? Well, it depends on what you're doing with your laptop. If you're using it for the Internet and light work, probably not much difference.

If you're working with more video-intensive applications, using a computer with shared memory might slow things down to an unacceptable level. This usually brings up two questions: What's the optimal balance? And where do I change this?

To answer the first question, again, it depends on what you are doing. If you perform more video-intensive operations (or if you're gaming), then you might want to set aside more memory for the video card. If you're not as concerned with rapid pixilation, then less is fine. Which brings us to the second question: Where do you set it? Shared memory is configured in the system BIOS/UEFI. Each BIOS/UEFI is different, so be sure to consult your owner's manual if you have any questions. Keep in mind that some BIOSs/UEFIs will allow you to set aside only a certain amount of memory—say, 512 MB to 2 GB—for video memory.

How does this affect your computer when you upgrade the memory? First, keep in mind that some of your memory will be taken by the video card, so you might want to upgrade to more memory than you originally had planned, if the laptop supports it. Second, after upgrading the memory, you will need to go into the BIOS/UEFI and reconfigure how much memory you want allocated to the video card.

Screen

The screen, as you would expect, shows the image just as it does on a full-sized, stand-alone display unit. It will function the exact same way as an LCD, LED, OLED, touchscreen, or whatever display technology you want to talk about. We won't review each of the different types of display technologies here—refer back to Volume 1 Chapter 3 if you need a refresher. What we will do, however, is provide a quick reminder of key components that make the screen work.

Backlight The *backlight* provides lighting for LCD screens. Recall that OLED screens don't need backlights.

Inverter For systems with a backlight, the *inverter* is needed to provide AC power to it. Inverters often look like a long, thin rod. Keep in mind that they can store energy, which can be shocking to the unwary technician.

Digitizer Many laptop displays are touchscreen, and the digitizer is the component that translates the touch into electrical signals for the computer. Most of the time on laptops, the glass part of the display contains the digitizer, webcam, and microphone.

Webcams are nearly universal on laptops today. The most common placement is right above the display on the laptop, although some are below the screen. Most laptops with webcams will also have a smaller circle to the left of the lens, which has a light that turns on to illuminate the user when the webcam is on. Some people are a bit paranoid about their webcams, and they will put something over the camera to block it when it's not in use. Many laptops will have a built-in, slidable shutter for this purpose.

Exam Objective 1.1, "Given a scenario, monitor mobile device hardware and use appropriate replacement techniques," highlights three components in the display you should know how to replace: Wi-Fi antenna connector, camera/webcam, and microphone.

If you're going to use a webcam to conduct a video call with someone, it helps if they can hear you too. That's where the *microphone* comes into play. Microphones are often built into the display unit as well. The webcam shown in Figure 4.7 has the illumination light to the left and the microphone inputs on both sides of the lens. Generally speaking, the illumination light will only be on when the webcam is active—it can be a good signal to the user as to whether they are "live" or not. Microphones can also be built into the bottom half of the clamshell, either above the keyboard or somewhere on the front bezel.

FIGURE 4.7 Webcam and microphone

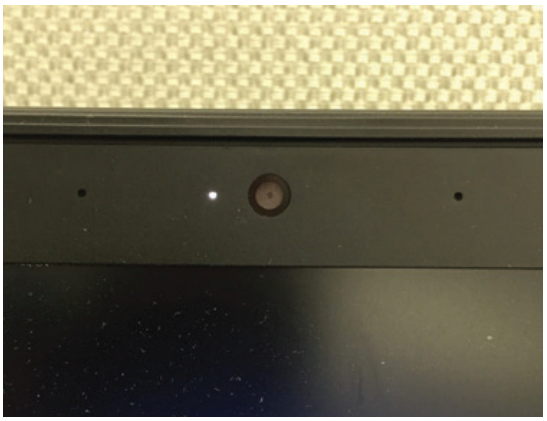

Wi-Fi Antenna

Practically all laptops produced today include built-in Wi-Fi capabilities. Considering how popular wireless networking is today, it only makes sense to include 802.11 functionality without needing to use an expansion card. With laptops that include built-in Wi-Fi, the *Wi-Fi antenna* is generally run through the upper half of the clamshell case. This is to get the antenna higher up and improve signal reception. The wiring will run down the side of the display and through the hinge of the laptop case and plug in somewhere on the motherboard.

The Wi-Fi antenna won't affect what you see on the screen, but if you start digging around in the display, know that you'll likely be messing with your wireless capabilities as well.

It's worth repeating again that the Wi-Fi antenna wiring usually runs through one of the hinges of the laptop case. Damage to a hinge can cause Wi-Fi problems. Later in the "Other Internal Components" section, we have an exercise on removing the Wi-Fi card, in which you will disconnect the antenna wires.

Replacing Components Within a Laptop Display

While going through the exercises in this chapter, you've probably noticed that the inside of a laptop is tight and complex, and navigating it can be a challenge. Laptop displays are the

same way. There are not as many components as compared to the bottom half of a laptop, but space is even tighter. Within a display, the components you might need to replace include the screen (or touchscreen/digitizer, which some manufacturers call the display panel), Wi-Fi antenna, webcam, microphone, and inverter.

In most cases, servicing components within the display means removing the display assembly from the base assembly. Exercise 4.2 shows you how to do that.

EXERCISE 4.2

Removing the Display Assembly

1. Remove the bottom cover.
2. Remove the wireless NIC from the base assembly.
3. Open the latch for the display cable and disconnect it from the motherboard. It's the left connector shown in Figure 4.8.

FIGURE 4.8 Display cable (left) and touchscreen board cable connectors (right)

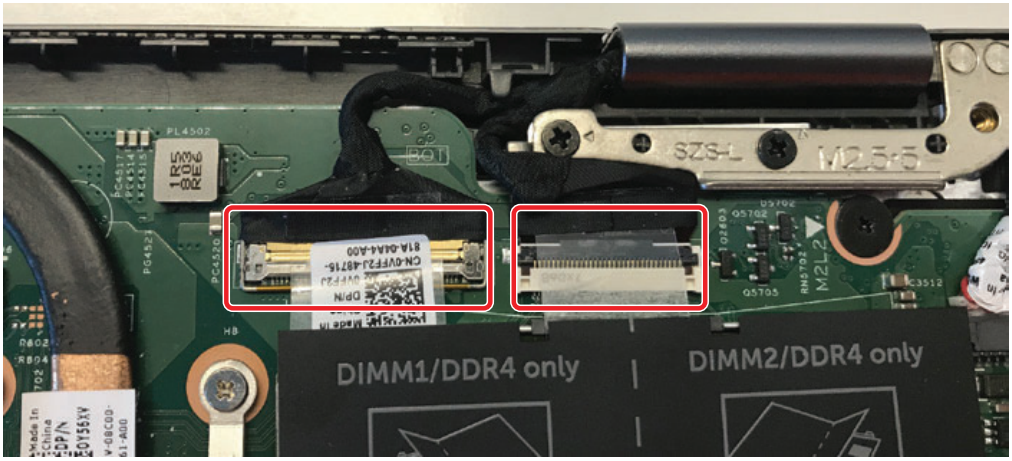

4. Disconnect the touchscreen board cable from the motherboard. It's the right connector shown in Figure 4.8.
5. Open the computer so that the display assembly is parallel with the base assembly, and lay it screen side down on a flat surface.

   We recommend that you place it on a soft surface to avoid scratching the screen.
6. Remove the screws that secure the display assembly to the base assembly. There are two screws for each hinge, and they are highlighted in Figure 4.9.

FIGURE 4.9 Display assembly hinge screws

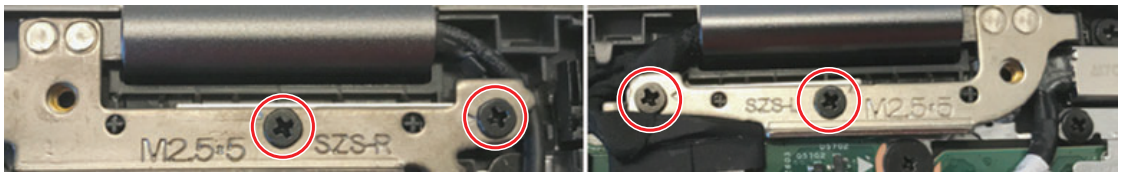

7. Lift the display assembly off of the base assembly.

   Note that the NIC antenna wires, display cable and connector, and touchscreen board cable and connector will go with the display assembly.

If you're replacing the entire display unit, all you need to do is get the new unit and reverse the steps you followed in Exercise 4.2. If you're replacing a component within the display unit, you need to go further. To get to the Wi-Fi antenna, webcam, microphone, and inverter, you must also remove the display panel. Exercise 4.3 gives you the general steps needed to accomplish this.

EXERCISE 4.3

Removing the Display Panel

1. Remove the bottom cover.
2. Remove the display assembly.
3. Remove the display cable from the routing guides inside the hinge covers.
4. Use a plastic scribe to pry apart the back of the display assembly and the display panel.

   Step 4 is a tricky step. There's not much of a gap between the display panel and the back of the display assembly. It's generally best to try to start prying it apart near the center of one of the sides of the panel—that's where you're likely to get the most separation to get the scribe in there.

   Also, we highly recommend using a plastic scribe. You can use a very thin flat screwdriver, but the metal is more likely to mark up or damage the case than the plastic scribe. Once you start getting separation, gently work your way around the edge of the display panel to separate the whole thing.

5. Identify the component you need to replace, and remove it.

EXERCISE 4.3 *(continued)*

Figure 4.10 shows a Lenovo ThinkPad with the display bezel removed. You can see wires running along the side and top. Some of the wires go to the webcam, and others are for the wireless antenna. As always, consult the laptop's documentation to identify the proper components.

FIGURE 4.10 Wires inside a laptop display

Motherboards and Processors

As with desktop computers, the motherboard of a laptop is the backbone structure to which all internal components connect. However, with a laptop, almost all components must be integrated onto the motherboard, including onboard circuitry for the USB, video, expansion, and network ports of the laptop. With desktop systems, the option remains to not integrate such components. Because of the similarities between laptop and desktop components, some material in the next few sections will be familiar to you if you have read Volume 1 Chapter 1, "Motherboards, Processors, and Memory."

Laptop Motherboards

The primary differences between a laptop motherboard and a desktop motherboard are the lack of standards and the much smaller form factor. As mentioned earlier, most motherboards are designed along with the laptop case so that all the components will fit inside. Therefore, the motherboard is nearly always proprietary, and that's what we mean by "lack of standards." They use the technologies you're used to, such as USB and Wi-Fi, but it's very unlikely that you're going to be able to swap a motherboard from one laptop to another, even if both laptops are from the same manufacturer. Figure 4.11 shows an example of a laptop motherboard. Its unusual shape is designed to help it fit into a specific style of case with the other necessary components.

To save space, components of the video circuitry (and possibly other circuits as well) may be placed on a thin circuit board that connects directly to the motherboard. This circuit board is often known as a *riser card* or a *daughterboard*; an example is shown in

Figure 4.12. We also labeled them in the Dell in Figure 4.4. They're harder to see when they're in the case with other components, but you may be able to tell that they have very different shapes from those in Figure 4.11 and Figure 4.12.

FIGURE 4.11 A laptop motherboard

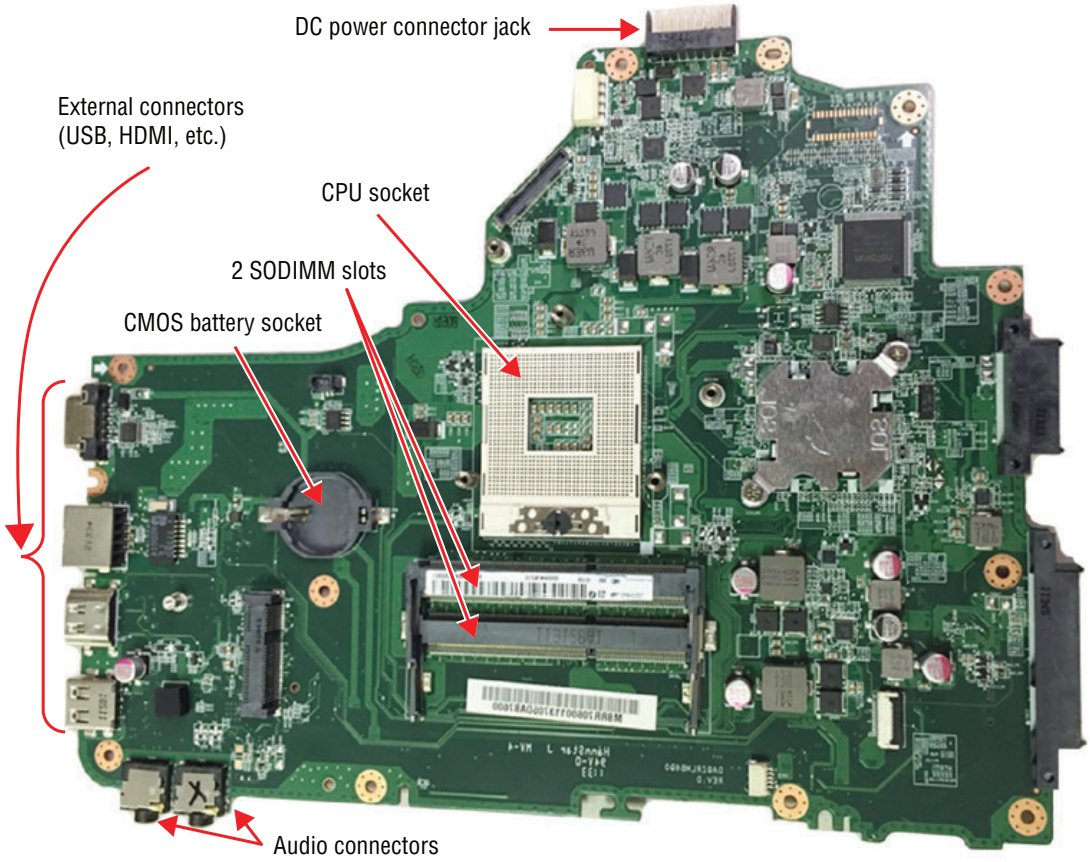

Dell happens to call this daughterboard an input/output board, or I/O board. This model also has another daughterboard. If you were ordering parts, you'd want to specify I/O board or daughterboard, but for general purposes, the term daughterboard is applicable.

Having components performing different functions (such as video, audio, and networking) integrated on the same board is a mixed bag. On one hand, it saves a lot of space. On the other hand, if one part goes bad, you have to replace the entire board, which is more expensive than just replacing one expansion card. Exercise 4.4 walks you through the steps to remove the motherboard from the Dell Inspiron 13 7000. As you'll see, you need to remove several components and disconnect several connectors before you can get the motherboard out.

FIGURE 4.12 A laptop daughterboard

EXERCISE 4.4

Removing the Motherboard from a Laptop

1. Remove the bottom cover.
2. Remove the battery.
3. Remove the fan.
4. Remove the heat sink.
5. Remove the solid-state drive.
6. Disconnect all the connectors plugged into the motherboard.

   Be careful, because two of them (for the I/O cable and the display cable) have a latch that needs to be released first. By our count, there are nine connectors. We've highlighted them in Figure 4.13. The two with latches have rectangles around them, and the rest have ovals.

FIGURE 4.13 Disconnecting the motherboard connectors

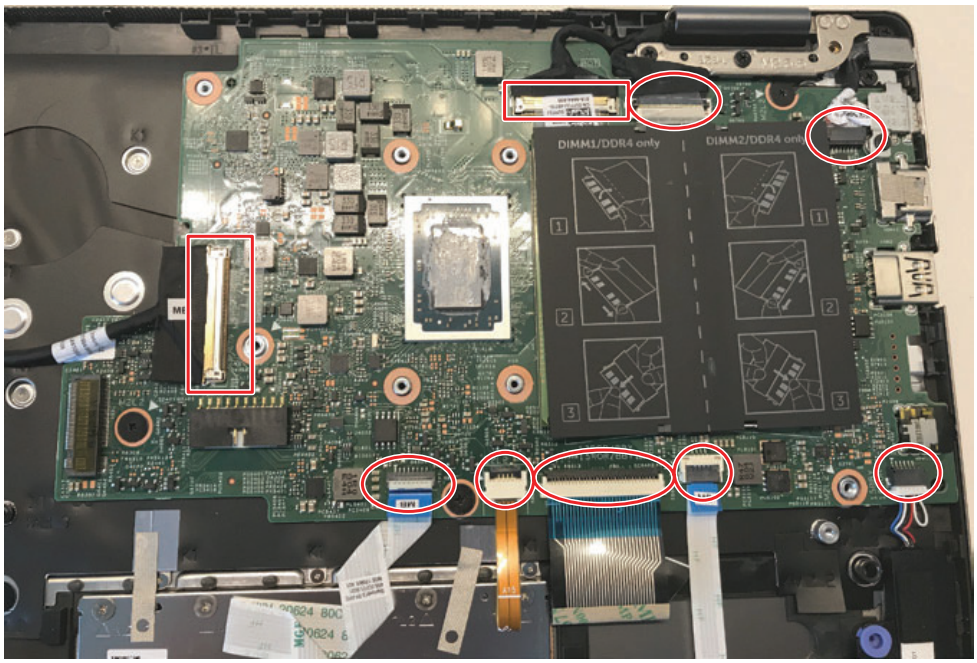

7. Remove the two screws from the system board bracket. They're shown in Figure 4.14 with circles around them.

FIGURE 4.14 Removing the screws

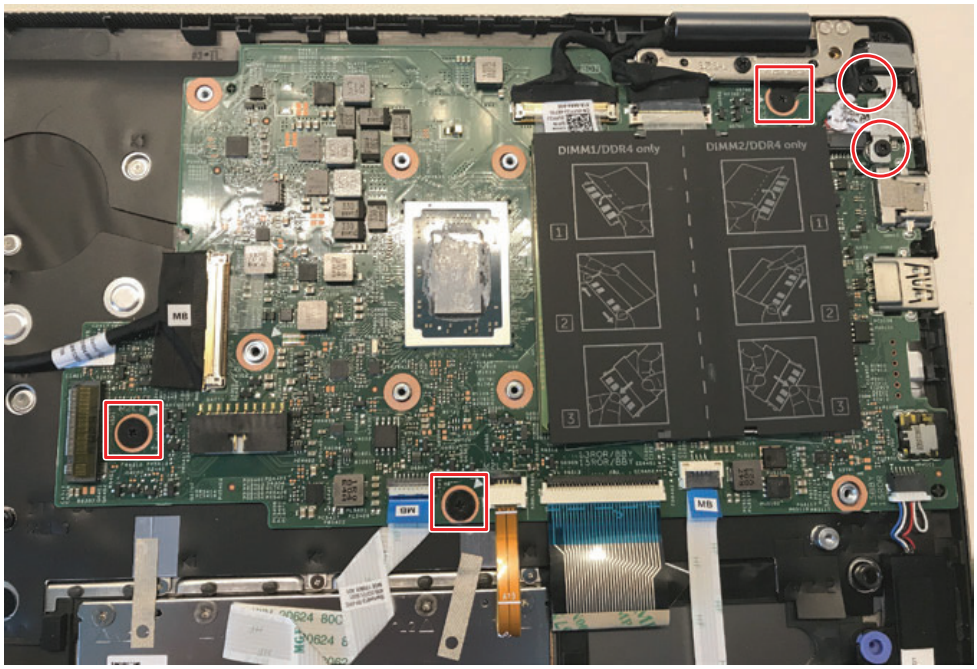

EXERCISE 4.4 *(continued)*

8. Remove the system board bracket.
9. Remove the three screws that hold the motherboard to the base assembly. In Figure 4.14, they are highlighted with squares around them.
10. Lift the motherboard off the base assembly.

To replace the motherboard, follow the steps in reverse order.

Laptop Processors

Just as with desktop computers, the central processing unit (CPU) is the brain of the laptop computer. And just like everything else, compared to desktop hardware devices, laptop hardware devices are smaller and not quite as powerful. The spread between the speed of and number of cores in a laptop CPU and that of a desktop model can be significant. Fortunately, the gap has closed over the years, and laptop processors today are pretty fast—most people think they perform just fine. It's up to the user to determine if the difference in speed hurts their usage experience.

Laptops have less space than desktops, and therefore the CPU is usually soldered onto the motherboard and is not upgradable. You can see the processor of the Dell we've been working on in Figure 4.14—it's the small silver square to the left of the RAM. Within confined computing spaces, heat is a major concern. Add to that the fact that the processor is the hottest-running component, and you can see where cooling can be an issue. To help combat this heat problem, laptop processors are engineered with the following features:

Streamlined Connection to the Motherboard Nearly all desktop processors mount using pin connectors, whether on the CPU or on the motherboard (as is the case with LGA sockets). Pins and sockets are big and bulky, meaning they're not a laptop's friends. Laptop processors are generally either soldered directly to the motherboard or attached using the Micro-FCBGA (Flip Chip Ball Grid Array) standard, which uses balls instead of pins. In most cases, this means that the processor cannot be removed, meaning no processor upgrades are possible.

Lower Voltages and Clock Speeds To combat heat, you can either slow the processor down (run it at a lower speed) or give it less juice (run it at a lower voltage). Again, performance will suffer compared to a desktop processor, but lowering heat is the goal here.

Active Sleep and Slowdown Modes Most laptops will run the processor in a lower power state when on battery power to extend the life of the battery. This is known as *processor throttling*. The motherboard works closely with the operating system to determine if the processor really needs to run at full speed. If it doesn't, it's slowed down to save energy and to reduce heat. When more processing power is needed, the CPU is throttled back up.

While many portable computers will have processors that have just as many features as their desktop counterparts, others will simply use stripped-down versions of desktop processors. Although there's nothing wrong with this, it makes sense that components specifically designed for laptops fit the application better than components that have been retrofitted for laptop use. Consider an analogy to the automobile industry: it's better to design a convertible from the ground up than simply to cut the top off an existing coupe or sedan.

Memory

Laptops don't use standard desktop computer memory chips, because they're too big. In fact, for most of the history of laptops, there were no standard types of memory chips. If you wanted to add memory to your laptop, you had to order it from the laptop manufacturer. Of course, because you could get memory from only one supplier, you got the privilege of paying a premium over and above a similar-sized desktop memory chip.

Exam Objective 1.1 includes RAM as a component you need to monitor and replace.

Fortunately, the industry seems to have settled on one form factor, which is the *small outline dual inline memory module (SODIMM)*, which we introduced in Volume 1 Chapter 1. Recall that they're much smaller than standard DIMMs, measuring 67.6 millimeters (2.6″) long and 32 millimeters (1.25″) tall. SODIMMs are available in a variety of configurations, including older 32-bit (72-pin and 100-pin) and 64-bit (144-pin SDRAM, 200-pin DDR/DDR2, 204-pin DDR3, 260-pin DDR4, and 262-pin DDR5) options. Different standards of DDR SODIMMs are only a few millimeters longer or shorter than other versions. You probably won't be able to tell the difference unless they are right next to each other, or unless you try to install them and they don't fit in the socket. (You should never have the latter problem, though, because you will check the documentation first!) Figure 4.15 shows a laptop DDR5 SODIMM above a desktop DDR5 DIMM for a size comparison.

FIGURE 4.15 Laptop SODIMM and Desktop DIMM

Laptop manufacturers may still opt to go the proprietary route due to design considerations that favor a custom solution. In fact, many ultra-thin laptop models will have the RAM soldered onto the motherboard to save space. In such systems, the memory is not upgradable. To see what kind of memory your laptop uses, check either the manual or the manufacturer's website. You can also check third-party memory producers' websites (such as `www.crucial.com`).

Just as with desktop computers, make sure the SODIMM you want to put into the laptop is compatible with the motherboard. The same standards that apply to desktop memory compatibility apply to laptops. This means that you can find DDR, DDR2, DDR3, DDR4, and DDR5 SODIMMs for laptops. DDR has topped out at 1 GB per module, whereas DDR2 and DDR3 SODIMM modules can be purchased in sizes up to 8 GB and DDR4 up to 32 GB and DDR5 up to 64 GB (at the time this book was being written). Exercise 4.5 shows you how to replace SODIMMs in a laptop.

EXERCISE 4.5

Replacing Laptop Memory

1. Remove the bottom cover.
2. Remove the battery.
3. Locate the system memory.

   In the Dell Inspiron 13 7000, the two RAM modules are under Mylar flaps.
4. Pull a Mylar flap back to reveal the RAM, as shown with one of the SODIMMs in Figure 4.16.

FIGURE 4.16 SODIMM in a laptop

5. Carefully spread apart the metal securing clips on each end of the memory module. It will pop up, as shown in Figure 4.17.

FIGURE 4.17 Removing a SODIMM

6. Lift the memory module out of the slot.

When replacing the SODIMM, be sure to align the notch, and insert the module at the approximate angle it was at when it popped up after being released. Slowly press down on the module until it clicks into place. If it doesn't click into place, it's not inserted properly.

For several years, a laptop memory form factor called MicroDIMM enjoyed some popularity. The MicroDIMM is over 50 percent smaller than a SODIMM—only about 45.5mm (about 1.75″) long and 30mm (about 1.2″) wide. Another major difference is that the MicroDIMM does not have any keying notches on the bottom. Figure 4.18 shows a 172-pin MicroDIMM compared to a DDR5 SODIMM; notice that it's a lot closer to being square-shaped than the SODIMM is. It was designed for the ultra-light and portable subnotebook style of computer. Common MicroDIMM form factors included 64-bit modules with 172 or 214 pins for DDR2/DDR3. DDR4 and DDR5 MicroDIMMs have not been produced. Replacing a MicroDIMM is just like replacing a SODIMM.

FIGURE 4.18 172-pin MicroDIMM and a SODIMM

Storage

Storage is important for every computer made. If you can't retrieve important files when you need them, the computer isn't very useful. While the trend is moving toward storing more data online (in the cloud), there's still considerable need for built-in storage.

Exam Objective 1.1 lists hard drives (HDDs) and solid-state drives (SSDs) as devices to monitor and know how to replace.

Laptop Hard Drives

Laptops don't have the room for the full-sized 3.5″ hard drives that desktop computers use. Smaller form factor drives at 2.5″ or 1.8″ that are less than ½″ thick are more appropriate. These drives share the same controller technologies as desktop computers; however, they use smaller connectors. Figure 4.19 shows an example of a standard 3.5″ hard drive compared to a 2.5″ laptop hard drive.

FIGURE 4.19 A 3.5″ hard drive (left) compared to a 2.5″ hard drive (right)

To save space and heat, almost all laptops today use a *solid-state drive (SSD)*, which we introduced in Volume 1 Chapter 2, "Expansion Cards, Storage Devices, and Power Supplies." Recall that, unlike conventional magnetic hard drives, which use spinning platters, SSDs have no moving parts. They use the same solid-state memory technology found in the other forms of flash memory. Otherwise, they perform just like a traditional magnetic HDD, except they're a lot faster.

Connecting a regular SSD in a desktop is usually just like connecting a regular HDD; they have the same Serial Advanced Technology Attachment (SATA) and power connectors. Laptops often have a specialized connector and a single cable that handles both data and power, as shown in Figure 4.20. Most manufacturers also make them in the same physical dimensions as traditional hard drives, even though they could be made much smaller, like removable flash drives. (This is probably to preserve the "look" of a hard drive so as to not confuse consumers or technicians.)

FIGURE 4.20 2.5″ SSD, motherboard connector, and cable

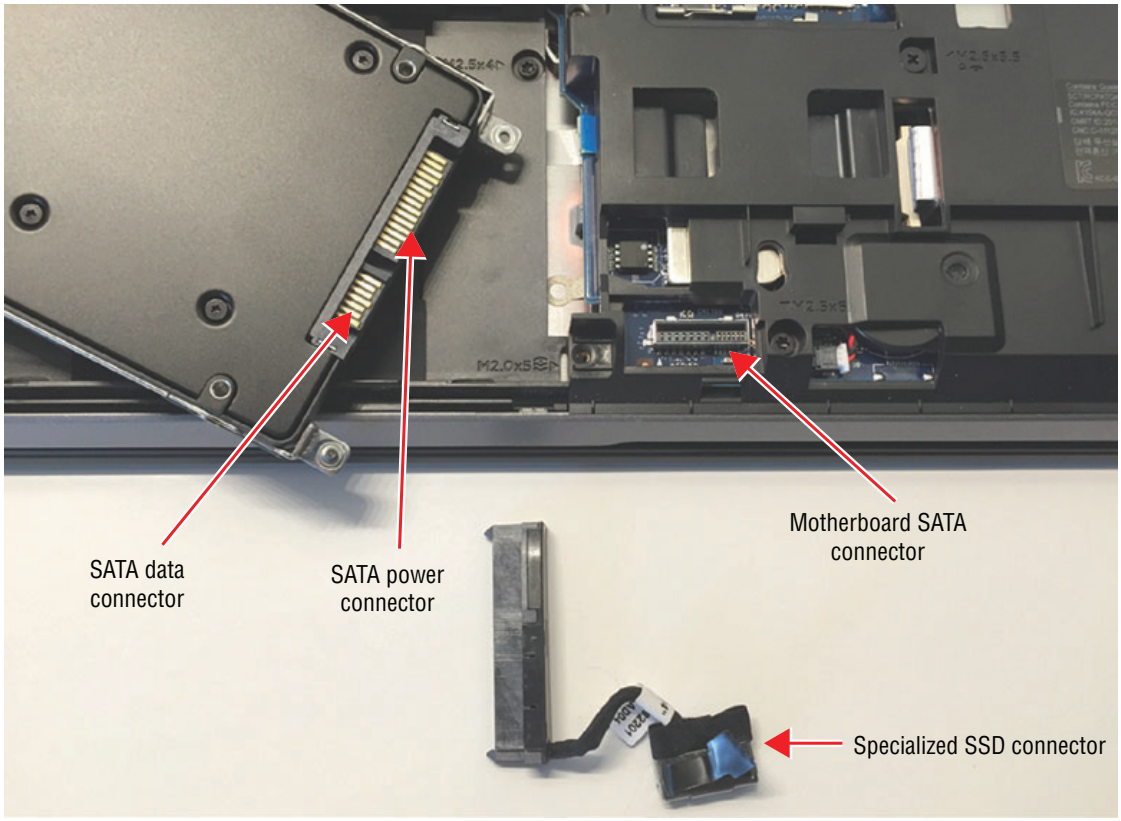

Newer SSDs may come in the even smaller M.2 form factor. In fact, the Dell Inspiron we've been working on in this chapter has an M.2 SSD. Figure 4.21 shows the SSD and the M.2 connector it plugs into. Exercise 4.6 walks you through removing an M.2 SSD from a laptop.

FIGURE 4.21 M.2 SSD and M.2 connector

EXERCISE 4.6

Removing an M.2 SSD from a Laptop

1. Remove the bottom cover.
2. Remove the screw that holds the SSD to the base assembly.

   The drive will pop up slightly. The screw is highlighted in Figure 4.22.
3. Slide the drive directly out of the connector.

FIGURE 4.22 Removing an M.2 SSD

Removing 2.5″ or 1.8″ SSDs from a laptop will require a few more steps than Exercise 4.6 did. First, disconnect the drive cable from the motherboard (refer to Figure 4.20), and then remove the two to four screws that hold the drive or its mounting bracket in place. Figure 4.23 shows an example of an SSD in a Lenovo ThinkPad that uses three screws—we've already removed them but highlighted the holes for you. Then lift out the drive (and bracket as needed). We show the disconnected drive next to the laptop in Figure 4.24.

FIGURE 4.23 2.5″ SATA SSD

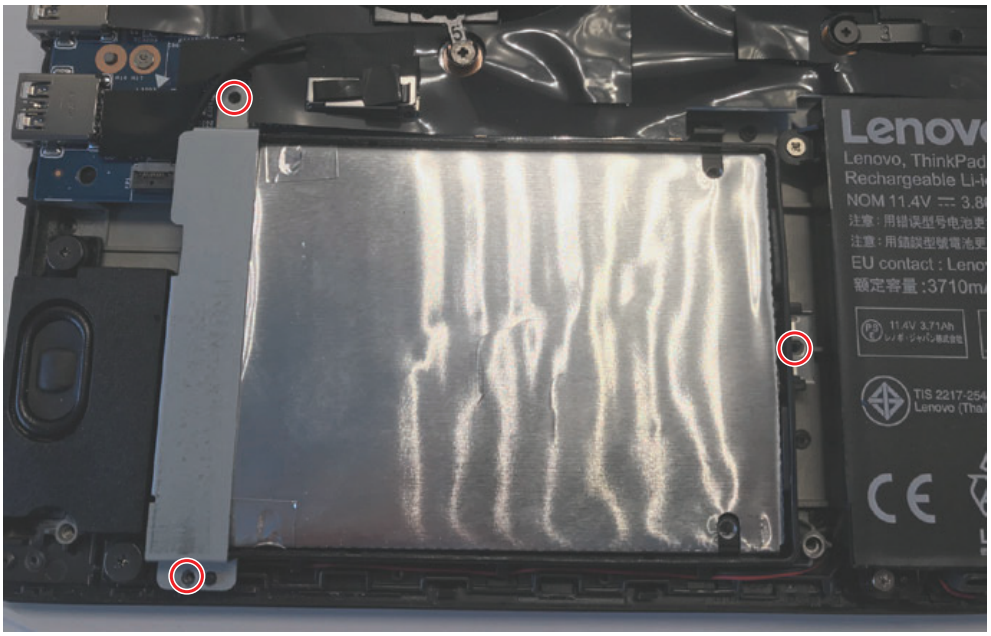

FIGURE 4.24 Disconnected SATA SSD

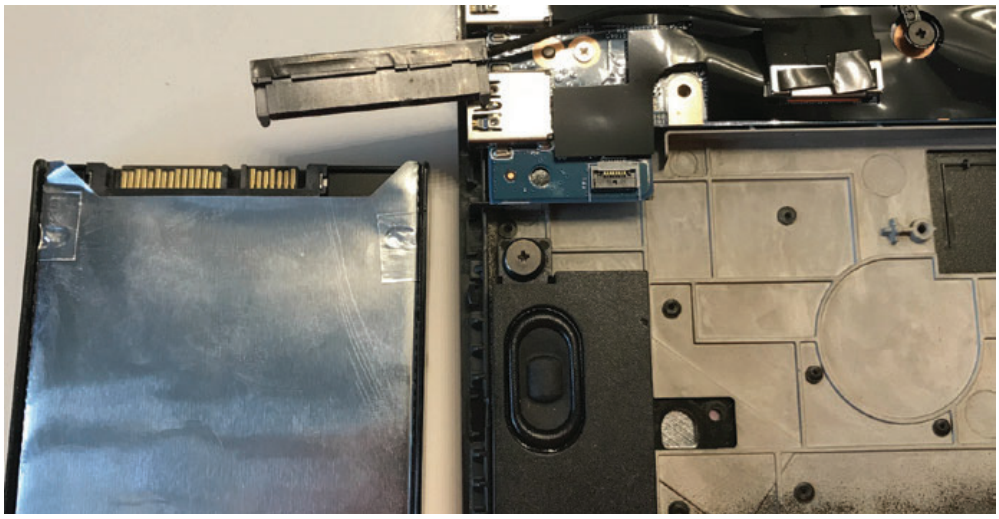

SSDs are perfect for laptops because they are faster, produce less heat, and are more shock resistant than conventional hard drives. They are more expensive per gigabyte, but that's the trade-off you have to make for better performance.

Hard Drive Data Migration

After replacing a hard drive or upgrading a laptop to a newer model, it may be necessary to get all of the user's data from the old drive (or laptop) onto the new one. Doing so is often referred to as data *migration* or hard drive migration. When setting up for a migration, there are two key questions to ask. The first is, what needs to be migrated? There's a big difference between moving someone's data and moving operating system settings and configurations. The second is, will the old drive be accessible when the new one is up and running? If so, there are more options to perform a migration. The answers to these will help determine which of the migration methods are preferable—either manual file copying or using migration software.

Copying Files Manually If all that's needed is to get user data from the old device to the new one, then manually copying files is usually an easy option. If the old drive will be inaccessible after the replacement, then files can be copied from the old drive to the cloud or to an external hard drive first. The drive replacement can be made, and then the files copied from the cloud or external drive to the new hard drive. If the old drive will still be accessible, then files can be copied across a network (if conditions permit) or copied from computer to computer with a transfer cable.

The downside to this method is that it doesn't transfer any settings or configurations (such as a user's desktop, color scheme, installed printers, etc.), and it normally doesn't work for apps. Apps will need to be reinstalled on the new drive, and the user will have to reconfigure their settings.

Using Migration Software Migration software can move files, settings, configurations, and apps from one drive to another. In most cases for this to work, both drives need to be accessible. For example, the old drive needs to be in a different operational computer, in a second expansion slot on the new computer, or otherwise connected, such as through a USB port using a USB-to-SATA adapter. If the connection is available, the migration software can work its magic.

There are several different migration apps available on the market. For example, Laplink PCmover (`https://go.laplink.com`) works with Windows computers. It has several versions, including the Home edition, which is good for normal users, and a Professional version that can copy multiple user profiles. Macrium Reflect X (`www.macrium.com`) is another option for Windows-based systems, and SuperDuper (`https://shirt-pocket.com`) works for macOS. All of these apps feature a graphical interface, where you choose what you want to migrate and the software takes it from there.

Laptop Optical Drives

Nearly all laptops have a hard drive, but rarely does a laptop made today have an internal optical drive. There just isn't room for one. If you need one, you can attach an external optical drive via an expansion port such as USB. It might be a bit slower than an internal drive, but it's better than not having one at all.

Input Devices

Because of the small size of laptops, getting data into them presents unique challenges to designers. They must design a keyboard that fits within the case of the laptop. They must also design some sort of pointing device that can be used with graphical interfaces like Windows. The primary challenge in both cases is to design these peripherals so that they fit within the design constraints of the laptop (low power and small form factor) while remaining usable.

Exam Objective 1.1 includes keyboard and keys as laptop hardware you should know how to replace.

Keyboards

A standard-sized desktop keyboard wasn't designed to be portable. It wouldn't fit well with the portable nature of a laptop. That usually means laptop keys are not normal size; they must be smaller and packed together more tightly. People who learned to type on a typewriter or a full-sized keyboard often have a difficult time adjusting to a laptop keyboard.

Keyboards may need to be replaced if keys are missing or are stuck and won't function. If a user has spilled a beverage onto their keyboard, it could cause problems that necessitate a replacement. On most laptop keyboards, you can't replace just one key; the whole board must be replaced to fix a malfunctioning key. If needed, an external USB keyboard can be added as a solution until the faulty built-in one is repaired.

Laptop keyboards are built into the lower portion of the clamshell. Sometimes, they can be removed easily to access peripherals below them, like memory and hard drives, as in the Lenovo ThinkPad series. Other times, removing the keyboard is one of the most challenging tasks, because nearly all other internal components need to be removed first. Exercise 4.7 illustrates the joy of removing the keyboard from the Dell Inspiron 13 7000.

EXERCISE 4.7

Removing a Laptop Keyboard

1. Remove the bottom cover.
2. Remove the battery.
3. Remove the SSD.
4. Remove the wireless card.

EXERCISE 4.7 *(continued)*

5. Remove the fan.
6. Remove the heat sink.
7. Remove the display assembly.
8. Remove the I/O board.
9. Remove the power adapter port.
10. Remove the motherboard.
11. Remove the speakers.
12. Remove the 14 screws that hold the keyboard shield to the base assembly.
13. Lift the keyboard from the base assembly.

As you can see, it was necessary to remove 11 components before we got to the keyboard. With a laptop like this, replacing the keyboard is not a quick-service item. This is also a great time to remember that you should always take pictures (if allowed) and document screw types and locations when tearing laptops apart. After removing 11 components, you might be dealing with several screw types and things can get confusing!

Pointing Devices

In addition to using the keyboard, you must have a method of controlling the onscreen pointer in the Windows (or other graphical) interface. Most laptops today include multiple USB ports for connecting a mouse, and you can choose from a wide range of wired or wireless full-sized or smaller mice. There are several additional methods for managing the Windows pointer. Here are some of the more common ones:

- Trackball
- Touchpad
- Point stick
- Touchscreen

Because of different pointing-device preferences, some laptops include multiple pointing devices to appeal to a wider variety of people.

Trackball Many early laptops used trackballs as pointing devices. A *trackball* is essentially the same as a mouse turned upside down. The onscreen pointer moves in the same direction and at the same speed that you move the trackball with your thumb or fingers.

Trackballs are cheap to produce. However, the primary problem with trackballs is that they do not last as long as other types of pointing devices; a trackball picks up dirt and oil from operators' fingers, and those substances clog the rollers on the trackball and prevent it from functioning properly.

Trackpad/Touchpad To overcome the problems of trackballs, a newer technology that has become known as the trackpad was developed. Many people call the trackpad a touchpad, but Touchpad is actually the trademarked name of Hewlett-Packard (HP) trackpad devices. They are one and the same, though.

A *trackpad* is a device that has a pad of touch-sensitive material. The user draws with their finger on the trackpad, and the onscreen pointer follows the finger motions. Included with the trackpad are two buttons for left- or right-clicking (although with some trackpads, you can perform the functions of the left-click by tapping on the trackpad, and Macs have one button). Figure 4.25 shows a trackpad on a laptop.

FIGURE 4.25 Laptop trackpad

Yet another name for trackpad-type devices is *drawing pad*. Typically, people only refer to it as a drawing pad if it's larger than a typical laptop trackpad and attached to a computer via USB. Drawing pads are often the size of a sheet of paper.

One problem people have with a trackpad is the location. You'll notice that the trackpad is conveniently placed right below the laptop keyboard, which happens to be where your palms rest when you type. Sometimes this will cause problems, because you can inadvertently cause your mouse cursor to do random things like jump across the screen. Most trackpads have settings to allow you to control the sensitivity, and they will also differentiate between a palm touching them and a finger. If you have a sensitive trackpad that is giving you trouble, you can disable it altogether. Exercise 4.8 shows you how to do that in Windows 11. The specific steps to disable the trackpad will differ by manufacturer—you will almost always be able to disable it through the operating system, and some laptops have a function key to disable it as well. The steps in Exercise 4.8 were performed on an HP ZBook laptop. Consult the laptop documentation if you are unable to locate the setting.

EXERCISE 4.8

Disabling a Touchpad in Windows 11

1. Open the Touchpad settings app by typing **touchpad** in the Windows search box and pressing Enter. You will see a screen similar to the one in Figure 4.26.

EXERCISE 4.8 *(continued)*

FIGURE 4.26 Touchpad settings

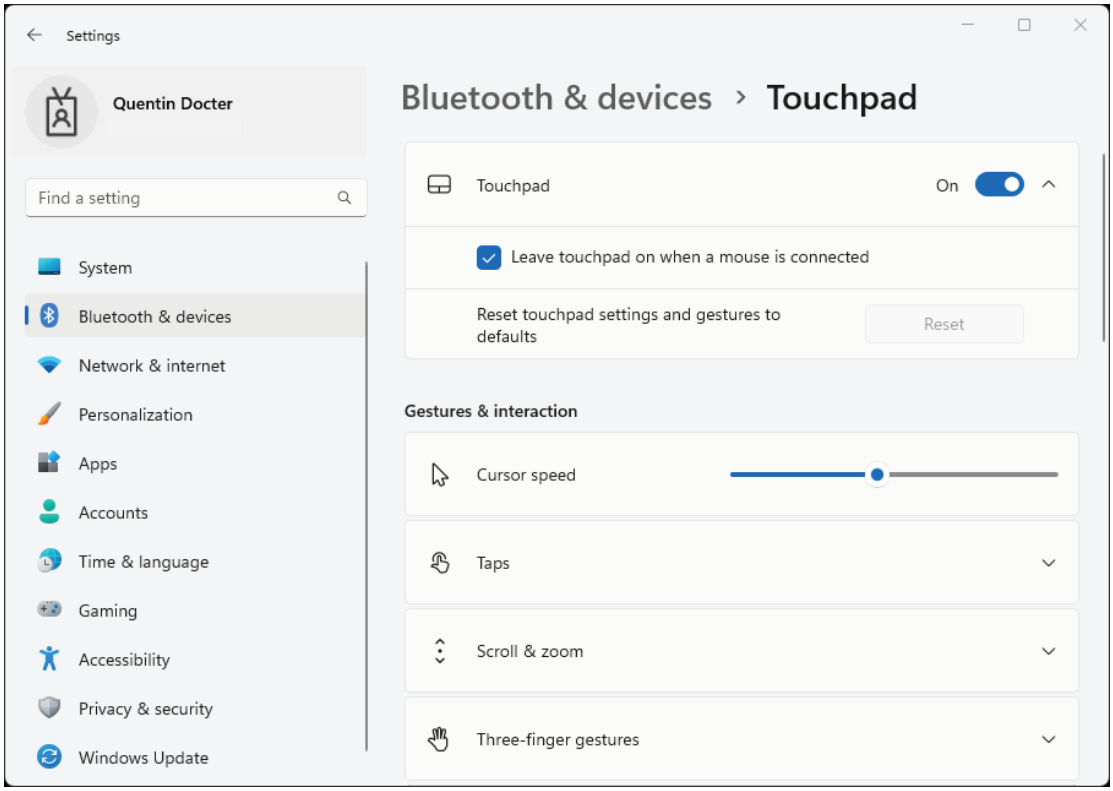

Remember that Touchpad is an HP trademark. On other laptops, trackpad settings are often in the Mouse Settings app.

2. Turn the Touchpad off by sliding the toggle bar to the off position. Notice there is also an option to leave the touchpad on (or turn it off) when a mouse is connected.
3. (Optional) Explore the other Touchpad settings, such as how to manage taps, scrolling and zooming, and multi-finger gestures.
4. Close the Touchpad window when finished.

The trackpad can be replaced if it becomes defective, just like most other laptop components. Before you replace it, first make sure that it's not just disabled in the operating system.

Track Point/Point Stick With the introduction of the ThinkPad series of laptops, IBM introduced a new feature known as the *Touchpoint*, generically known as a *point stick* or a *track point*. (ThinkPads are now produced by Lenovo.) The track point is a pointing device that uses a small rubber-tipped stick. When you push the track point in a particular direction, the onscreen pointer goes in the same direction. The harder you push, the faster the onscreen pointer moves. The point allows fingertip control of the onscreen pointer, without the reliability problems associated with trackballs. Figure 4.27 shows a track point.

FIGURE 4.27 Track point

Track points can have their own problems, however. The stick may not return to center properly, causing the pointer to drift when not in use. You might also notice the rubber cover for the stick becoming a bit gummy with extended use. Most manufacturers supply replacement covers of varying textures with new systems. Exercise 4.9 shows you how to disable and manage track points.

EXERCISE 4.9

Disabling a Track Point in Windows 11

1. Open the Mouse settings app by typing **`mouse`** in the Windows search bar and pressing Enter.

   You may be able to type **`pointstick`** in the search bar and press Enter to open Mouse Settings as well.

2. In Mouse Settings, scroll down to the PointStick Settings area and click menu arrow on the right side. You will get a Properties screen like the one in Figure 4.28.

EXERCISE 4.9 *(continued)*

FIGURE 4.28 Point stick properties

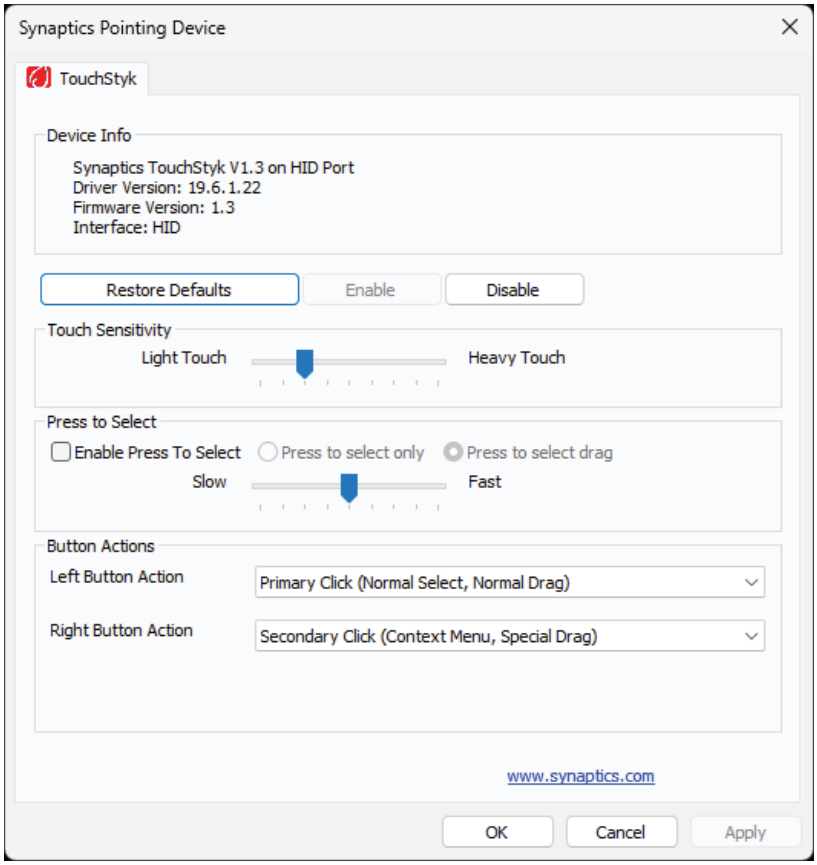

3. Click Disable to disable the track point.
4. (Optional) Explore other configuration options for the track point.
5. Once you're satisfied with the configuration, click OK to close the Properties window.

Touchscreen *Touchscreens* are standard fare for smartphones and tablets and many laptop computers as well. We've already introduced the technology as a display device, so here we'll cover it as an input device. The idea is pretty simple: it looks like any other display device, but you can touch the screen and the system senses it. It can be as simple as registering a click, like a mouse, or it can be more advanced, such as capturing handwriting and saving it as a digital note.

Although the technical details of how touchscreens work are beyond the scope of this chapter, there are a few things you should know:

- Some touchscreens work with any object touching them, whereas others require a conductive input, such as your finger. iPhones are a great example of this, to

which anyone who lives in cold climates and wears gloves can attest.

- Some touchscreens are coated with a film that is sensitive to touch. Cleaning these screens with regular glass cleaner can ruin the touchscreen nature of the device. It's best to clean those devices only with a damp cloth, as needed.
- Some touchscreens need to be calibrated in order to properly interpret the user's input. See the manufacturer's documentation for details.

User options to configure touchscreens are often limited to calibrating the screen for touch input or pen input. If the touchscreen seems to be sensing input incorrectly or not detecting input from the edges of the screen, it might be time to recalibrate it. In Windows, open Control Panel ➢ Tablet PC Settings, and then click Calibrate. Windows will ask the user to specify pen or touch input. Once that's specified, the user will be able to draw on the screen with their desired input tool, and the system will sense it. (In most cases, there will be a crosshairs on the screen and the user needs to touch where it is.) At the end, the user needs to save the calibration data for it to take effect.

Many laptops sold today are advertised as 2-in-1 devices, meaning that the top portion of the clamshell that holds the display folds completely over, effectively turning the laptop into a glorified tablet. These devices use touchscreens, and when they are in tablet mode, the keyboard is disabled. Some will also automatically change the screen orientation to landscape or portrait, depending on how it's being held, just like a smartphone or tablet.

Internal Expansion

Although laptop computers are less expandable than their desktop counterparts, many can be expanded to some extent. The two primary forms of internal expansion used in laptops today are Mini PCIe and M.2.

Exam Objective 1.1 lists wireless cards as devices you should know how to replace in a laptop. To be safe, you should understand how to replace both Mini PCIe and M.2 expansion cards.

Mini PCIe

Since around 2005, *Mini PCIe* has been the most common slot for laptop expansion cards. We introduced PCIe in Volume 1 Chapter 2, and mini PCIe is just like the full version except the connectors are smaller. These cards reside inside the case of the laptop and are connected via a 52-pin card edge connector. Mini PCIe cards come in two sizes. The full-sized cards are 30mm wide and 51mm long. Half-sized cards (one is shown in Figure 4.29, with the connector at the bottom) are 30mm wide and 27mm long. Mini PCIe cards support USB and PCIe x1 functionality, and at the same speeds. Additionally, Mini PCIe cards have the 1.5V and 3.3V power options.

FIGURE 4.29 Mini PCIe card in a laptop

M.2

We also introduced M.2 in Volume 1 Chapter 2, so we won't go into a great amount of detail here. Figure 4.21 earlier in this chapter shows an M.2 hard drive and slot. For purposes of comparing it to Mini PCIe, know that it uses a narrower connector (22mm vs. 30mm) that has more pins (66-pin vs. 52-pin). M.2 supports USB 2.0 and newer. The slowest M.2 slots support PCIe x2, and M-keyed slots support PCIe x4, making it much faster than Mini PCIe. Most M.2 expansion cards focus on communications or storage. Common types of cards you will see in the market include the following:

- Wi-Fi network cards
- Bluetooth cards
- Cellular cards
- SSDs

Many laptops don't come with any free internal expansion slots, either M.2 or Mini PCIe. For example, the Dell we've been working on has two M.2 slots, but both are filled. One has

the SSD, and the other is used by the Wi-Fi card. Be sure to check the documentation before you buy an expansion card for a laptop, to see what it supports.

Batteries and Power Adapters

Because portable computers have unique characteristics as a result of their portability, they have unique power systems as well. Portable computers can use either of two power sources: batteries or adapted power from an AC or DC source. Regardless of the source of their power, laptops utilize DC power to energize their internal components. Therefore, any AC power source needs to be rectified (converted) to DC. Most laptop display backlights, on the other hand, require high-voltage, low-amperage AC power. To avoid a separate external AC input, an inverter is used to convert the DC power that is supplied for the rest of the system to AC for the backlight. In case it's not obvious, converters and inverters perform opposite functions, more or less.

Understand how to replace batteries in a laptop, per Exam Objective 1.1.

Batteries

There are many different battery chemistries that come in various sizes and shapes. Nickel cadmium (NiCd), lithium-ion (Li-ion), and nickel-metal hydride (NiMH) have been the most popular chemistries for laptop batteries. A newer battery chemistry, lithium-polymer (Li-poly), has been gaining in prominence over recent years for smaller devices.

Laptop batteries may be easily removable without cracking the case. Figure 4.30 shows a removable Li-ion battery for an HP laptop that just pops off the back side of the case.

FIGURE 4.30 A removable laptop Li-ion battery

A removable battery is very easy to replace in the event of a battery failure. However, most laptops today make use of an internal battery, such as the one shown in Figure 4.31. This particular battery is very thin—only 5mm (⅕″) thick. Exercise 4.10 shows you how to remove an internal laptop battery.

EXERCISE 4.10

Removing an Internal Laptop Battery

If a laptop won't power on while on battery power but it works while plugged in and the battery shows it has a charge, a simple re-seating of the battery might solve the problem. Unplug it, let the laptop sit for a few seconds, and then plug it back in.

1. Remove the bottom cover.
2. Disconnect the battery from the motherboard.

   Its connector is highlighted with a rectangle in Figure 4.31.
3. Remove the four screws holding the battery in place.

   In Figure 4.31, the screws are highlighted with circles.
4. Lift the battery out of the base assembly.

FIGURE 4.31 Internal laptop battery

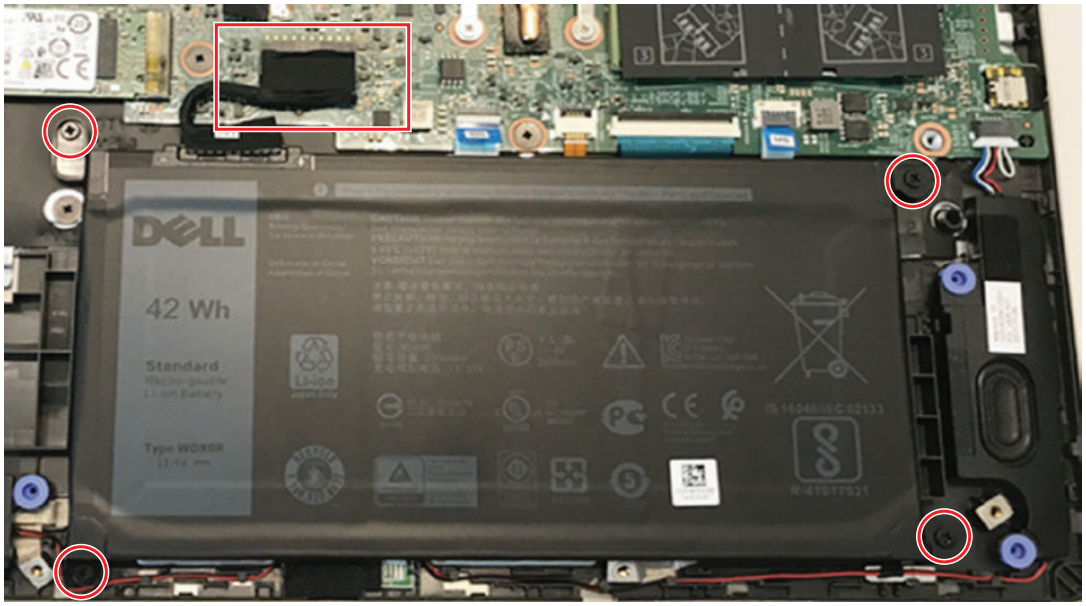

Even though replacing an internal battery requires slightly more effort than replacing an external battery, it's still generally a simple operation.

Battery chemistries can be compared by energy density and power density. Energy density measures how much energy a battery can hold. Power density measures how quickly the stored energy can be accessed, focusing on access in bursts, not prolonged runtime. An analogy to the storage and distribution of liquids might help solidify these concepts. A gallon bucket has a higher "energy density" and "power density" than a pint bottle; the bucket holds more and can pour its contents more quickly. Another common metric for battery comparison is rate of self-discharge, or how fast an unused battery reduces its stored charge.

Real World Scenario

Is That Battery Really Dead?

Some batteries, such as nickel cadmium (NiCd) batteries, suffer from a performance-affecting chemical memory loss. Others, such as lithium-ion (Li-ion), don't suffer from this affliction but do suffer from so-called digital memory loss that plagues the built-in gauges that monitor the charge left in the battery. This effect can be observed in software gauges that read the battery's charge level. The digital memory effect manifests itself as a sudden loss of power when the gauges register, say, 30 percent remaining capacity. The fix, much like the fix for chemical memory in NiCd batteries, is to allow a full discharge once a month or so. This is called *battery calibration* and can be performed right in the device while it's using the battery.

Other than this occasional full discharge, Li-ion batteries last longer when you partially discharge them and then recharge them, making them ideal for laptops and personal handheld devices, such as smartphones, that tend to get used sporadically on battery power before being plugged back in to charge.

Power Adapters

Most laptop computers can also use AC power with a special adapter (called an *AC adapter*) that converts AC power input to DC output. The adapter can be integrated into the laptop, but more often it's a separate "brick" with two cords: one that plugs into the back of the laptop and another that plugs into a wall outlet. Figure 4.32 is a photo of the latter. Most AC adapters use a USB-C connection, which has been popular since 2014. Some laptop manufacturers will still roll out proprietary connectors, though. In those cases, only an AC adapter that matches the connector will work.

Use caution when selecting a replacement AC adapter for your laptop. You should choose one rated for the same or higher wattage than the original. Regarding the input voltage of the adapter, care must also be taken to match the adapter to the power grid of the surrounding region. Most adapters are autoswitching (universal voltage), but some adapters may have a fixed AC input requirement. Purchasing the wrong unit can result in a lack of functionality or damage to the laptop. If or when working internationally, you should always ascertain whether some sort of converter is required, even for autoswitching adapters.

FIGURE 4.32 A laptop AC adapter

Other Internal Components

There are a few internal components we've referenced in chapter exercises, but we haven't given explicit details on how to remove them. That's what this section is for. Note that these components aren't currently on the exam objectives, but it helps to know how to remove them anyway. The four components we'll examine are the fan, heat sink, wireless NIC, and CMOS battery. And since we'll be talking about CMOS, which is related to the BIOS/UEFI, we'll look at how to upgrade (or flash) the BIOS/UEFI as well. Exercise 4.11 shows you how to remove the system fan.

EXERCISE 4.11

Removing the System Fan

1. Remove the bottom cover.
2. Remove the two screws that hold the system fan in place. They are highlighted in Figure 4.33.
3. Lift the fan away from the base assembly.

FIGURE 4.33 Removing the system fan

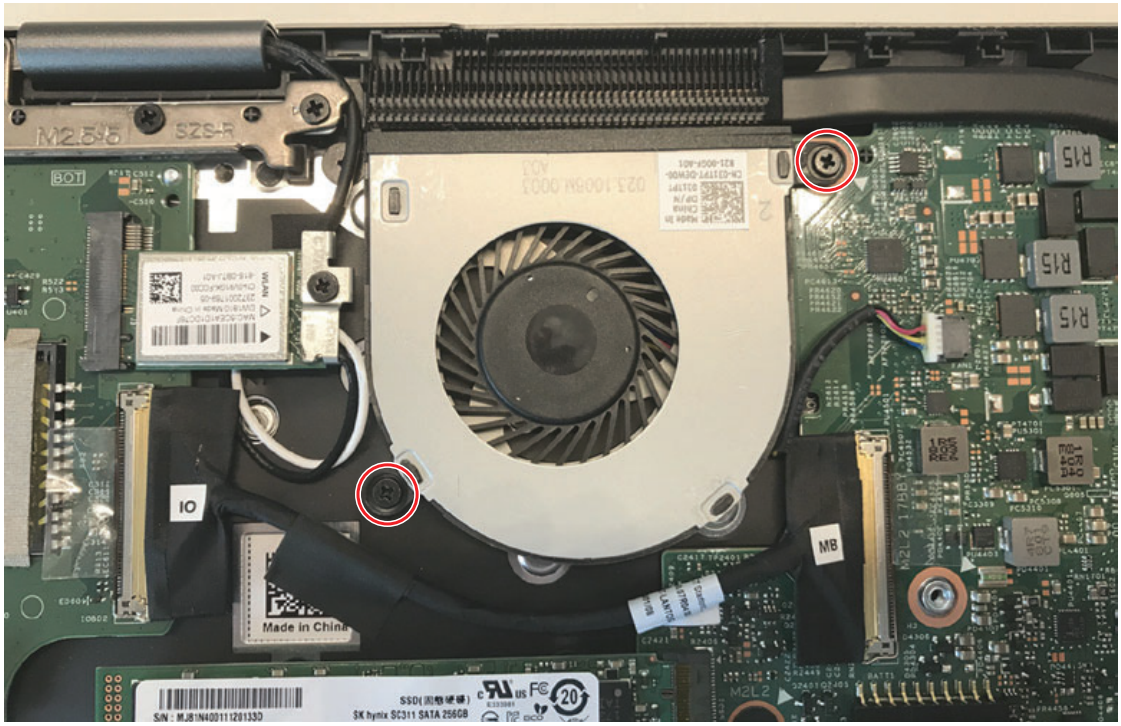

Now that the fan is removed, Exercise 4.12 shows you how to remove the CPU heat sink.

EXERCISE 4.12

Removing the CPU Heat Sink

1. Remove the bottom cover.
2. Remove the system fan.
3. Loosen the four screws holding the heat sink in place, as highlighted in Figure 4.34.

   The screws will not completely come out of the heat sink bracket.
4. Lift the heat sink straight up off the processor.

   It might stick a bit, since most are adhered with a thermal compound.

EXERCISE 4.12 *(continued)*

FIGURE 4.34 Removing the heat sink

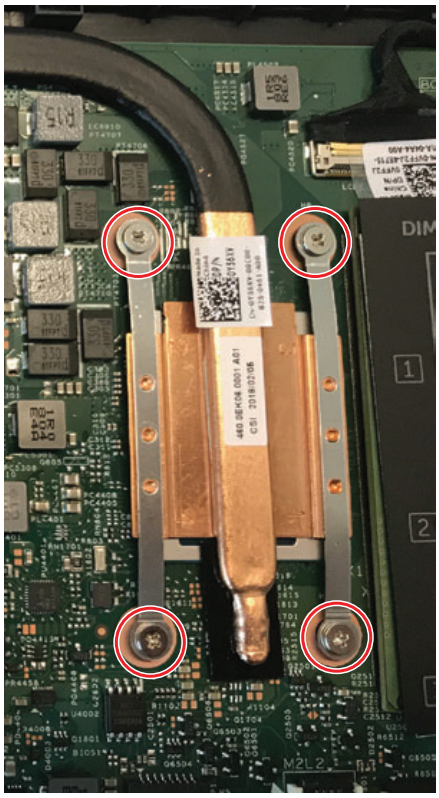

Exercise 4.13 shows you how to remove the wireless NIC. Perhaps ironically, it actually does have wires, but those are for the antenna that runs up into the display.

EXERCISE 4.13

Removing the Wireless NIC

1. Remove the bottom cover.
2. Remove the screw holding the wireless card bracket in place. It's circled in Figure 4.35.

FIGURE 4.35 Removing the wireless NIC

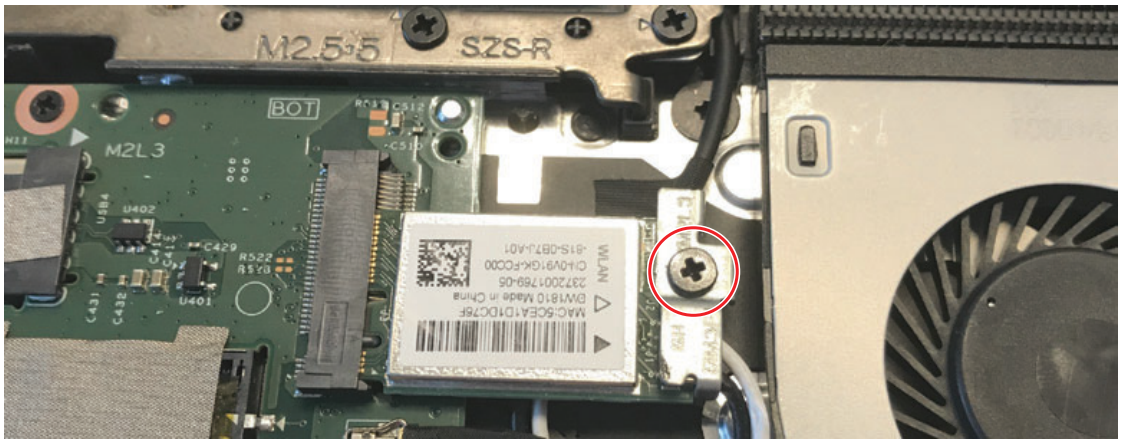

3. The wireless card will pop up, as shown in Figure 4.36.

FIGURE 4.36 Disconnecting the two antenna wires

4. Disconnect the two antenna wires that are attached to the end of the card.
5. Pull the wireless NIC straight out of the M.2 socket.

When you're reconnecting the wireless card, the white antenna cable will go on the main post, which is indicated by a white triangle. The black antenna cable attaches to the auxiliary connector, which is marked with a black triangle.

Exercise 4.14 shows how to remove the CMOS battery. In most laptops, the CMOS battery is covered by a black rubber coating.

EXERCISE 4.14

Removing the CMOS Battery

1. Remove the bottom cover.
2. Disconnect the connector from the motherboard. It's highlighted with a box in Figure 4.37.

FIGURE 4.37 Disconnect the CMOS battery

3. Peel the battery off the base assembly.

   Odds are good that it's been glued into place, so you might have to pull hard.

Replacement batteries often come with a small amount of adhesive to secure the new battery in place.

Flashing the system BIOS/UEFI is usually a pretty straightforward process. You can download a BIOS/UEFI update from the manufacturer's website and then run the program. Exercise 4.15 shows sample steps for flashing the BIOS on a Dell laptop.

EXERCISE 4.15

Flashing the System BIOS

1. Turn on the computer.
2. Visit `https://www.dell.com/support`.
3. Navigate to Drivers & Downloads.
4. Select the operating system installed on your laptop.
5. Find the BIOS section.
6. Select the file and click Download to download the BIOS.
7. After the download is complete, locate the file on your laptop.
8. Double-click the file to run it and begin the BIOS update process. Follow the onscreen instructions.

Physical Privacy and Security Components

Mobile devices are mostly self-contained units, which aids in their portability. One downside to the portability is that mobile devices can easily be carried away by someone other than their rightful owner. And, if a user is doing work on a laptop in a public place, others could shoulder-surf and see things they're not supposed to. The use of physical privacy and security components can help deter these unwanted behaviors.

Exam Objective 1.1 directs you to know about two physical privacy and security components: biometrics and near-field scanners.

Biometrics

Many mobile devices allow you to log in or unlock the screen through *biometrics*, or the use of a body part. For example, you can unlock your smartphone by enabling the facial recognition feature, or some older models have a fingerprint scanner for the same purpose. Other options for higher-end security systems include voice recognition and retinal scanning.

Laptops may have a fingerprint scanner built into the keyboard, such as the square one shown to the right of the Intel sticker in Figure 4.38. Others may be a rectangle or a circle. If a laptop doesn't have a biometric scanner and you want to add one, there are many USB options available.

The use of biometrics can increase device security. Someone may be able to guess your password or see you type it in and can hack you that way. But fingerprints and other biometric features are unique. You may have seen movies where a super-secret spy ring replicates someone's fingerprint to gain access to a system, but in real life that type of thing is virtually unheard of.

FIGURE 4.38 Laptop fingerprint scanner

Near-Field Scanner Features

In Volume 1 Chapter 9, "Wireless and SOHO Networking," we will cover the technical details of *near-field communication (NFC)*. For now, know that NFC is used extensively for mobile payment systems because it's fast and convenient. The payment system is enabled through the use of a *near-field scanner*.

NFC is a wireless communications method, so the signals sent to and from NFC devices could be intercepted by a malicious third device. Because the maximum distance for NFC is about 10 centimeters (about 4″), it would be pretty difficult for someone to intercept NFC data without the sender and receiver knowing it, but it's not impossible. Confidentiality, denial of service (DoS) and on-path attacks are potential security concerns with NFC. The key thing to remember is, if using NFC for payment, be aware of your surroundings and keep an eye out for potentially suspicious-looking electronic devices.

Cable Locks

One way that you can help to physically secure your laptop is through the use of a physical laptop lock, also known as a *cable lock*. Essentially, a cable lock anchors your device to a physical structure, making it nearly impossible for someone to walk off with it. Figure 4.39 shows a cable with a keyed lock. With others, a small combination lock is used. If you grew up using a bicycle lock, these will look really familiar.

FIGURE 4.39 A cable lock

Sometimes a cable lock is called a *Kensington lock* (named after a company that makes them) or a *K-lock*.

Here's how it works. First, find a secure structure, such as the permanent metal supports of your workstation at work. Then, wrap the lock cord around the structure, putting the lock through the loop at the other end. Finally, secure the lock into your cable lock hole on the back or side of your laptop (Figure 4.40), and you're secure. If you forget your combination or lose your key, you're most likely going to have to cut through the cord, which will require a large cable cutter or a hacksaw.

FIGURE 4.40 Cable lock in a laptop

If someone wants your laptop bad enough, they can break the case and dislodge your lock. Having the lock in place will deter most people looking to make off with it, though.

Privacy Screens

To keep curious and prying eyes from viewing a laptop (or desktop) monitor, users can install a privacy screen over the front of their display. It's a thin sheet of semi-transparent plastic that reduces the viewing angle so that a display can be read only by someone directly in front of it. An example made by `TechStar.com` is shown in Figure 4.41.

FIGURE 4.41 Laptop privacy screen

Some laptops will have a built-in privacy screen setting that can be activated by the function (Fn) keys. For example, it would be Fn+F2 on the laptop keyboard shown in Figure 4.42. In our experience, the built-in privacy screens aren't quite as effective at cutting down the viewing angle as add-on privacy screens are, but they are better than nothing.

FIGURE 4.42 Fn+F2 enables the privacy screen.

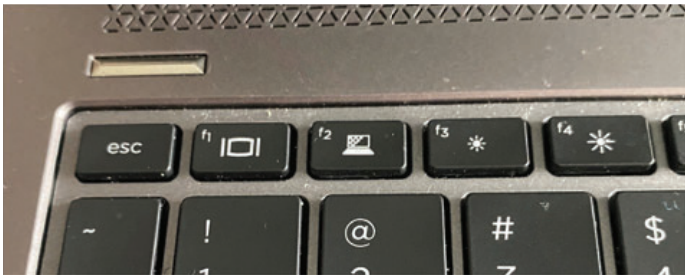

Setting Up and Configuring Accessories and Ports

Other than needing to plug in and charge for a while, laptops and mobile devices don't require anything to be plugged into them to provide a user with full functionality. With that said, there are a number of accessories that can be used to enhance functionality. These include touch pens, headsets, speakers, webcams, trackpads/drawing pads, docking stations, and port replicators. We'll examine each one of them after we take a quick look at accessory connection methods.

Connection Methods

Smaller mobile devices such as smartphones and tablets generally don't have very many physical ports on them. One or two at the most is all you're going to get. That of course limits expandability unless you are connecting the accessory using wireless methods. Laptops often give you more physical connectors—you're likely to have at least a few USB ports to play with, along with an audio jack and the power port.

Exam Objective 1.2, "Compare and contrast accessories and connectivity options for mobile devices," lists five connection methods you should be familiar with. They are USB and its variants, Lightning, NFC, Bluetooth, and tethering/hotspot.

We've covered all the connection methods listed in the A+ Exam Objective 1.2 for Exam 220-1201 in previous chapters with the exception of tethering/hotspot. Because of that, here we'll just do a quick refresh of the ones you need to know.

USB USB absolutely lives up to its name of being universal. There are multiple USB types to know, including microUSB, miniUSB, and the oval-shaped USB-C. Just about any type of accessory you could possibly want will come with a USB adapter.

Lightning Proprietary to Apple devices, the Lightning connector was the only external connector you would find on iPhones and iPads for several years. Macs and MacBooks might or might not have Lightning ports as well—if they do, they are intended mostly for device charging. Apple started moving away from Lightning in favor of USB-C in 2022, but expect Lightning to still be common for several years.

Near-Field Communication NFC is primarily used for mobile commerce and applications like scanning RFID tags, but it could be used for accessories as well. It's more likely to be used *by* an accessory, though, such as a mobile commerce scanner, than it is to connect one. Of the wireless technologies commonly used with mobile devices (Wi-Fi, Bluetooth, NFC), it has the shortest range (about 10 centimeters, or 4″).

Bluetooth As far as mobile accessories go, Bluetooth is a very popular connection technology. Mobile headsets, speakers, keyboards, and mice are all commonly connected to mobile devices using Bluetooth. Do keep in mind that if you have enabled Bluetooth on a small device such as a smartphone, it will eat up battery life a lot faster.

Tethering/Hotspot A mobile *hotspot* is when a device shares its cellular Internet connection with Wi-Fi capable devices. For example, a Wi-Fi-enabled laptop could use Wi-Fi to connect to a smartphone hotspot and get Internet access through the smartphone's cellular Internet connection. This tends to be relatively slow Internet, given cellular connection speeds, but it's better than nothing.

Tethering is when you have connected a device to a mobile hotspot. The term used to be reserved only for when you connected via USB cable, as opposed to connecting via wireless. Some devices will not function as a mobile hotspot but will allow you to tether a laptop (or other device) to it, so that the mobile device can share the cellular Internet connection.

Mobile Accessories

We could spend an entire chapter talking about mobile device accessories. There are dozens of mobile accessories on the market, including security devices, input/output tools, communication enablers, and mobile commerce endpoints. Here we will focus on the ones you need to know for the A+ exam. Most of them have already been covered in this book, so this will be a quick review of those and an introduction to the stylus.

Mobile accessories you need to be familiar with, per Exam Objective 1.2, are stylus, headset, speakers, and webcam. Docking stations, port replicators, and trackpad/drawing pad/track points are also listed in the same objective.

Stylus A *touch pen*, also known as a *stylus*, is a pen-shaped accessory used to write with or as a pointer. Touch pens come in a variety of shapes and sizes, although most have either a narrow tip (like a pen) or a soft rubber ball-like tip, kind of like a pencil eraser. The idea is that a touch pen will act as an input device on a touchscreen, enabling

freeform writing, drawing, or clicking through answers to a questionnaire, such as at a medical office. Several touch pens are shown in Figure 4.43.

FIGURE 4.43 Several touch pens

Source: Craig Spurrier / Wikimedia Commons / CC BY 2.5

Headset A *headset* is a common mobile device that delivers both output and input. The output is in the form of speakers, either in-ear or over-ear, and input comes from the attached microphone. Headsets come in a huge range of options, both wired and wireless (usually Bluetooth), from cheap models for conference calls to high-end, audiophile-grade gems. Figure 4.44 shows an inexpensive USB headset.

FIGURE 4.44 USB headset

Speakers *Speakers* are common audio output devices. Like headsets, a wide variety of speakers are available. USB and Bluetooth are the two most common options.

Webcam *Webcams* are video input devices. Most mobile devices have a built-in webcam. Stand-alone USB and Bluetooth ones are available as well.

Trackpad/Drawing Pad/Track Point We covered these accessories earlier in the chapter in the "Input Devices" section. Recall that these are input devices. Most of the time, trackpads and track points are built into laptop keyboards. Drawing pads can be added as an accessory for those who need larger freeform drawing spaces. Some have buttons, like the one shown in Figure 4.45, to provide additional functionality such as on-pad menu management or erase and undo features.

FIGURE 4.45 Drawing pad accessory with stylus

Source: Selftaken / Wikimedia Commons / CC BY SA 3.0

Docking Stations and Port Replicators

While users love laptops for their portability, sometimes a user simply needs to work with a full-sized keyboard or monitor. In addition, some users might have other external peripherals they need to plug in, such as external storage or a printer. It can be a pain to plug and unplug these devices constantly.

A *docking station* (shown in Figure 4.46) is a laptop-only peripheral that is basically an extension of the motherboard of a laptop. Because a docking station is designed to stay behind when the laptop is removed, it can contain things like a full-sized drive bay and expansion slots. The docking station in Figure 4.46 has the following ports: four USB (two are on the side), two DisplayPort, one HDMI, one RJ45, and a cable lock. It connects to the PC through a USB-C connection. It's about 5″ square and 2″ high.

FIGURE 4.46 A USB-C docking station

The docking station functions as a *port replicator*. A port replicator reproduces the functions of the ports on the back of a laptop so that peripherals such as monitors, keyboards, printers, and so on that don't travel with the laptop can remain connected to the dock and don't all have to be unplugged physically each time the laptop is taken away.

What's the Difference?

The difference between a docking station and a port replicator today is miniscule, because they basically provide the same functionality. Technically speaking, the only real difference is in how they connect to the laptop. (Some docking stations will contain built-in bays for storage or optical drives, but those are becoming increasingly rare.)

Older laptops will have a connector on the bottom to connect to a specific type of docking port, making docking ports proprietary. These ports work only with docking stations designed by the laptop's manufacturer, and vice versa. An example from the bottom of an HP laptop is shown in Figure 4.47. The HP docking station would have a corresponding connector.

FIGURE 4.47 HP laptop docking port

Today, almost all laptop port-expanding devices connect via USB-C, typically in the same port that would have been otherwise used for the power cord. These are technically port replicators, and they are not proprietary to one manufacturer.

In the real world, the terms get used interchangeably. Some manufacturers will call USB-C devices docking stations. It really doesn't matter which term is used, as long as the end user gets the functionality they need.

Summary

In this chapter, you learned about laptop and mobile device hardware. We discussed differences between laptops, mobile devices, and desktops, including the various components that make up smaller devices and how they differ in appearance and function from those on a desktop. We also talked about principles for disassembling and reassembling smaller devices in order to not lose parts.

We broke the chapter into several sections to discuss different components. We started with the case, and then moved on to displays, motherboards and processors, memory, storage, input devices, internal expansion, batteries and power adapters, physical privacy and security, other internal components, and external peripherals. In each section, we covered specifics related to laptop versions of these components and included exercises to show you how to service them.

The chapter puts particular emphasis on how to replace various components in the system, which is an important skill for many service technicians.

Finally, we ended the chapter by examining accessories and their connection methods. Connection methods included USB, Lightning, serial, NFC, Bluetooth, and tethering/hotspot. Accessories to remember include touch pens/stylus, trackpads/drawing pads/track points, headsets, speakers, webcams, and docking stations and port replicators.

Exam Essentials

Understand how to install and configure laptop components. Components include the battery, keyboard, RAM, hard drives and solid-state drives, and wireless cards.

Know laptop physical privacy and security components. They include biometrics and near-field scanner features.

Understand the components that make up a display in a mobile device. Display components include the screen, which may be a touchscreen or include a digitizer; Wi-Fi antenna; camera or webcam; microphone; and inverter.

Know how to connect mobile device accessories. Connection methods include USB, Lightning, serial, near-field communication (NFC), Bluetooth, and tethering/hotspot.

Be familiar with mobile device accessories. Ones to know include touch pens, trackpads/drawing pads/track points, stylus, headsets, speakers, webcams, docking stations, and port replicators.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. A client has a laptop with which you are unfamiliar. You are asked to perform a memory upgrade. How can you obtain the service manual for this laptop computer?
   - **A.** By pressing F1 while in Windows
   - **B.** By pressing F2 while the system is booting up
   - **C.** By reading the paper copy that comes with the laptop
   - **D.** By searching the manufacturer's website
2. Your manager tells you to replace the inverter on a broken laptop. Where will you find this component?
   - **A.** Underneath the keyboard
   - **B.** In the display
   - **C.** Next to the heat sink and fan
   - **D.** Near the CPU
3. You need to replace an LCD on a defective laptop. Which of the following are components of an LCD? (Choose all that apply.)
   - **A.** Inverter
   - **B.** Screen
   - **C.** Wi-Fi card
   - **D.** Backdrop
4. A user complains that his mouse cursor randomly jumps to different places on the screen when he types. Which device might you want to disable to prevent this from happening?
   - **A.** Trackpad
   - **B.** Mouse
   - **C.** Point stick
   - **D.** Trackball
5. A user needs to get their laptop connected to the Internet, but no Wi-Fi connections are available. They do have their smartphone with them. What is the best option for them to get Internet access for their laptop?
   - **A.** NFC
   - **B.** Tethering
   - **C.** Lightning
   - **D.** Bluetooth

6. You need to replace defective DDR4 memory in a laptop. How many pins will the replacement memory module have?
   - **A.** 200
   - **B.** 204
   - **C.** 260
   - **D.** 288

7. A user complains that his 2-in-1 laptop screen does not properly detect where he places his finger on the screen. What should you do first to resolve the issue?
   - **A.** Replace the touchscreen.
   - **B.** Replace the trackpad.
   - **C.** Recalibrate the touchscreen.
   - **D.** Recalibrate the trackpad.
   - **E.** Degauss the touchscreen.
   - **F.** Degauss the trackpad.

8. You need to wirelessly connect two accessories to a smartphone. Which of the following could be valid options to connect these wireless accessories? (Choose two.)
   - **A.** USB
   - **B.** Lightning
   - **C.** NFC
   - **D.** Bluetooth

9. There has recently been a string of hardware thefts in your office building. What should you recommend that your users implement to help avoid this problem?
   - **A.** Enable GPS tracking.
   - **B.** Use a docking station.
   - **C.** Use a cable lock.
   - **D.** Use the screensaver lock.

10. Which of the following describes the job of an inverter in a laptop?
    - **A.** It changes the screen orientation when the laptop is rotated.
    - **B.** It converts AC power into DC power for the display backlight.
    - **C.** It converts DC power into AC power for the display backlight.
    - **D.** It allows the laptop display to function in tablet mode.

11. A user is getting a new laptop. Which of the following is *not* a valid option to get his files and data from the old hard drive to the new hard drive?
    - **A.** Manually copy files from the old drive to the cloud, and then to the new drive.
    - **B.** Manually copy files from the old drive to an external hard drive, and then to the new drive.
    - **C.** Use a hard drive migration app.
    - **D.** Use the NFC disk transfer utility.

12. A user on your network wants to be able to draw images on a screen and have the laptop capture them. What type of device needs to be installed on their laptop?

A. Inverter

B. Capturer

C. Digitizer

D. Trackpad

13. A user has requested a laptop memory upgrade. The technician attempts to put in new memory, but the memory module is barely too wide for the slot. What is the most likely scenario?

A. The technician is trying to install DDR4 into a motherboard with DDR3 slots.

B. The technician is trying to install DDR3 into a motherboard with DDR4 slots.

C. The technician is trying to install MicroDIMMs into SODIMM slots.

D. The technician is trying to install DIMMs into SODIMM slots.

14. You have a user who needs to keep desktop devices such as keyboard, monitor, and mouse permanently connected so that they can be used by an attached laptop. What type of device do you recommend that they use?

A. Docking station

B. Keyboard, video, mouse (KVM) switch

C. Input/output virtual server

D. USB hub

15. The process by which the processor slows down to conserve power is officially called _____.

A. Underclocking

B. Cooling

C. Disengaging

D. Throttling

16. You need to replace a failed AC adapter for a client's laptop. When replacing the adapter, which of the following purchases is acceptable to obtain the same or better results?

A. An AC adapter with a higher voltage rating than the original

B. An AC adapter with a higher wattage rating than the original

C. A DC adapter with the same voltage rating as the original

D. An AC adapter with a lower voltage and wattage rating than the original

17. What should you do for a Li-ion battery that appears to charge fully but does not last as long as the battery's meter indicates that it will last?
    - **A.** Replace the battery.
    - **B.** Exercise the battery.
    - **C.** Calibrate the battery.
    - **D.** Short the terminals to discharge the battery.

18. You are giving a seminar to new technicians on replacing laptop hard drives. A student asks how do laptop hard drives differ from desktop hard drives. What do you tell them?
    - **A.** Laptop hard drives use completely different standards from those used by desktop hard drives for communication with the host.
    - **B.** Laptop hard drives are solid-state, whereas desktop hard drives have spinning platters.
    - **C.** Laptop hard drives require a separate power connection, whereas desktop hard drives are powered through the drive interface.
    - **D.** The most common form factors for laptop hard drives is smaller than that of a desktop hard drive.

19. One of your network users has an older laptop with no Bluetooth support. He recently received a Bluetooth headset that he needs to use with his laptop. What is the quickest and cheapest way to make the headset compatible with his laptop?
    - **A.** Buy a Wi-Fi–compatible headset.
    - **B.** Buy a USB Bluetooth adapter.
    - **C.** Buy a Mini PCIe Bluetooth adapter.
    - **D.** Replace the laptop.

20. A user wants her laptop to have a full-sized optical drive and external hard drive. Which of the following accessories is most likely to have full-sized bays for these devices?
    - **A.** Laptop base
    - **B.** Port replicator
    - **C.** Docking station
    - **D.** Mini PCIe dock

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answer compares to the authors', refer to Appendix B.

The hard drive on a Dell Inspiron 13 7000 computer failed. You have an extra hard drive of the exact same type. What would you do to replace it?

Chapter

5

Printers and Multifunction Devices

THE FOLLOWING COMPTIA A+ 220-1201 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

✓ **3.0 Hardware**

✓ **3.7 Given a scenario, deploy and configure multifunction devices/printers and settings.**

- Properly unbox device and consider set-up location
- Use appropriate drivers for a given operating system
  - Printer Control Language (PCL) vs. postscript
- Firmware
- Device connectivity
  - USB
  - Ethernet
  - Wireless
- Public/shared devices
  - Printer share
  - Print server
- Configuration settings
  - Duplex
  - Orientation
  - Tray settings
  - Quality
- Security
  - User authentication
  - Badging

        - Audit logs
        - Secured prints
    - Network scan services
        - Email
        - SMB
        - Cloud services
    - Automatic document feeder (ADF)/flatbed scanner

✓ **3.8 Given a scenario, perform appropriate printer maintenance.**

- Laser
    - Maintenance: Replace toner, apply maintenance kit, calibrate, and clean
- Inkjet
    - Ink cartridge, printhead, roller, and feeder
    - Maintenance: Clean printheads, replace cartridges, calibrate, and clear jams
- Thermal
    - Feed assembly
    - Special thermal paper
    - Maintenance: Replace paper, clean heating element, and remove debris
- Impact
    - Multipart paper
    - Maintenance: Replace ribbon, printhead, and paper

Even as technology progresses and almost all of our lives seem digitized, our society is still reliant on paper. When we conduct business, we use different types of paper documents, such as contracts, letters, and, of course, money. And because most of those documents are created on computers, printers are inherently important. Even with electronic business being the norm in many situations, you will likely still have daily situations that require an old-fashioned hard copy of something.

Printers are electromechanical output devices that are used to put information from the computer onto paper. They have been around since the introduction of the computer. Other than the display monitor, the printer is the most popular output device purchased for a computer because a lot of people want and sometimes need to have paper copies of the documents they create.

In this chapter, we will discuss the details of each major type of printing technology, including impact printers, inkjet printers, laser printers, and thermal printers. Once we cover the different types, we'll talk about installing and configuring printers and finish up with a section on printer maintenance.

Printer troubleshooting is covered in Volume 1 Chapter 13, "Mobile Device, Display, and Printer Troubleshooting."

Understanding Print Technologies and Imaging Processes

Several types of printers are available on the market today. As with all other computer components, there have been significant advancements in printer technology over the years. Most of the time, when faced with the decision of purchasing a printer, you're going to be weighing performance versus cost. Some of the higher-quality technologies, such as color laser printing, are relatively expensive for the home user. Other technologies are less expensive but don't provide the same level of quality.

In the following sections, you will learn about the various types of print technologies that you will see as a technician as well as their basic components and how they function. Specifically, we are going to look at four classifications of printing—impact, inkjet, laser, and thermal.

Impact Printers

The most basic type of printer is in the category known as an *impact printer*. Impact printers, as their name suggests, use some form of impact and an inked *printer ribbon* to make an imprint on the paper. Impact printers also use a paper feed mechanism called a *tractor feed* that requires special paper. Perhaps you've seen it before—it's continuous-feed paper with holes running down both edges.

There are two major types of impact printers: daisy-wheel and dot-matrix. Each type has its own service and maintenance issues.

For Exam Objective 3.8, "Given a scenario, perform appropriate printer maintenance," you need to be familiar with multipart paper as it relates to impact printers. We will cover more maintenance concepts in the "Performing Appropriate Printer Maintenance" section later in this chapter.

Daisy-Wheel Printers

The first type of impact printer to know about is the *daisy-wheel printer*. This is one of the oldest printing technologies in use. These impact printers contain a wheel (called the daisy wheel because it looks like a daisy) with raised letters and symbols on each "petal" (see Figure 5.1). When the printer needs to print a character, it sends a signal to the mechanism that contains the wheel. This mechanism is called the *printhead*. The printhead rotates the daisy wheel until the required character is in place. An electromechanical hammer (called a *solenoid)* then strikes the back of the petal containing the character. The character pushes up against an inked ribbon that ultimately strikes the paper, making the impression of the requested character.

FIGURE 5.1 A daisy-wheel printer mechanism

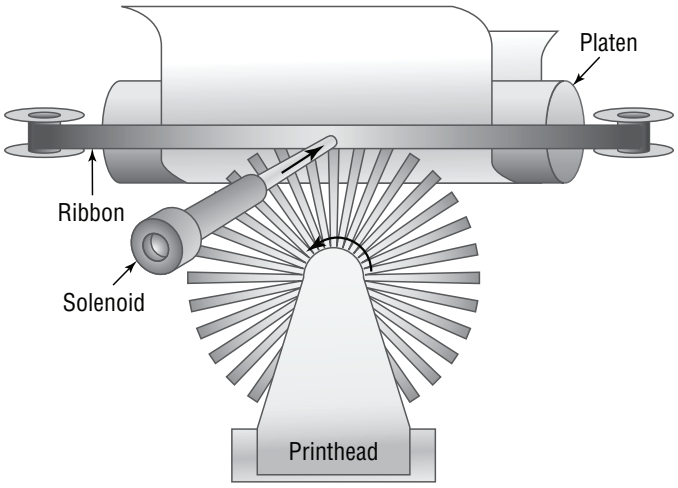

Daisy-wheel printers were among the first types of impact printer developed. Their speed is rated by the number of *characters per second (cps)* they can print. The earliest printers could print only two to four characters per second. Aside from their poor speed, the main disadvantage of this type of printer is that it makes a lot of noise when printing—so much so, in fact, that special enclosures were developed to contain the noise. There is also no concept of using multiple fonts; the font is whatever the character on the wheel looks like.

The daisy-wheel printer has a few advantages, of course. First, because it is an impact printer, you can print on multipart forms (like carbonless receipts), assuming that they can be fed into the printer properly. Sometimes, you will hear this type of paper referred to as *multipart paper* or *impact paper*. Second, it is relatively inexpensive compared to the price of a laser printer of the same vintage. Finally, the print quality is easily readable; the level of quality was given a name: *letter quality (LQ)*. Today, LQ might refer to quality that's better than an old-school typewriter (if you're familiar with them) but not up to inkjet standards.

Dot-Matrix Printers

The other type of impact printer to understand is the *dot-matrix printer*. These printers work in a manner similar to daisy-wheel printers, but instead of a spinning, character-imprinted wheel, the printhead contains a row of pins (short, sturdy stalks of hard wire). These pins are triggered in patterns that form letters and numbers as the printhead moves across the paper (see Figure 5.2).

FIGURE 5.2 Formation of images in a dot-matrix printer

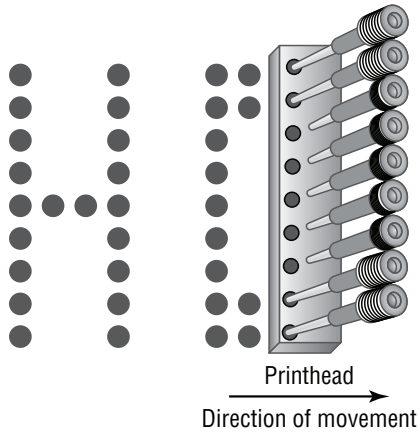

The pins in the printhead are wrapped with coils of wire to create a solenoid and are held in the rest position by a combination of a small magnet and a spring. To trigger a particular pin, the printer controller sends a signal to the printhead, which energizes the wires around the appropriate print wire. This turns the print wire into an electromagnet, which repels the print pin, forcing it against the ink ribbon and making a dot on the paper. The arrangement of the dots in columns and rows creates the letters and numbers that you see on the page. Figure 5.2 illustrates this process.

The main disadvantage of dot-matrix printers is their image quality, which can be quite poor compared to the quality produced with a daisy wheel. Dot-matrix printers use patterns of dots to make letters and images, and the early dot-matrix printers used only nine pins to make those patterns. The output quality of such printers is referred to as *draft quality*—good mainly for providing your initial text to a correspondent or reviser. Each letter looked fuzzy because the dots were spaced as far as they could be and still be perceived as a letter or image. As more pins were crammed into the printhead (17-pin and 24-pin models were eventually developed), the quality increased because the dots were closer together. Dot-matrix technology ultimately improved to the point that a letter printed on a dot-matrix printer was *almost* indistinguishable from daisy-wheel output. This level of quality is known as *near letter quality (NLQ)*.

Most impact printers have an option to adjust how close the printhead rests from the ribbon. So if your printing is too light, you may be able to adjust the printhead closer to the ribbon. If it's too dark or you get smeared printing, you may be able to move the printhead back.

Dot-matrix printers are noisy, but the print wires and printhead are covered by a plastic dust cover, making them quieter than daisy-wheel printers. They also use a more efficient printing technology, so the print speed is faster (typically starting around 72 cps). Some dot-matrix printers (like the Epson DFX series, which can run up to 1550 cps) can print close to a page per second! Finally, because dot-matrix printers are also impact printers, they can use multipart forms. Because of these advantages, dot-matrix printers quickly made daisy-wheel printers obsolete.

Inkjet Printers

One of the most popular types of printers in use today is the *inkjet printer*. As opposed to impact printers, which strike the page, these printers spray ink on the page to form the image. Inkjet printers typically use a reservoir of ink, a pump, and a nozzle to accomplish this. Older inkjet printers were messy, noisy, and inefficient, but the technology is good enough now that you see plenty of photo printers using inkjet technology. You might also hear these types of printers referred to as bubble-jet printers, but that term is copyrighted by Canon. You can think of inkjets as spraying droplets of ink in a very high-definition dot-matrix pattern, although printer manufacturers would likely scoff at the comparison to an older technology.

In the following sections, you will learn the parts of an inkjet printer as well as how inkjet printers work.

Parts of a Typical Inkjet Printer

Inkjet printers are simple devices. They contain very few parts (even fewer than dot-matrix printers) and, as such, are inexpensive to manufacture. It's common today to have a $40 to $50 inkjet printer with print quality that rivals that of basic laser printers.

Inkjet parts you need to know for Exam Objective 3.8 are the ink cartridge, printhead, roller, and feeder. We will cover a few additional components as well. They aren't specifically called out in the exam objectives, but we cover them for two reasons. One, so you understand how the entire printing process works and can troubleshoot problems in real life. Two, you will be tested on troubleshooting printers, and it helps to understand the process from start to finish to know where the problem is occurring.

The printer parts can be divided into the following categories:

- Printhead/ink cartridge
- Head carriage, belt, and stepper motor
- Paper feed mechanism
- Control, interface, and power circuitry

Why So Cheap?

The cost of inkjet printers is pretty low—one would think that the printer companies make very little money on them (which is true). So why do they price them so low?

The answer is that printer companies make the vast majority of their profits from ink cartridge sales. They figure that if they can get you into their printer, you will buy their ink refills—which is true—so you are then a captive consumer. Ink cartridges are very cheap to produce and have a high profit margin. Industry experts estimate that it costs about $2 to make an ink cartridge. On average, black cartridges cost about $20–$50, and color cartridges are about $20–$60.

Printhead/Ink Cartridge

The first part of an inkjet printer is the one that people see the most: the *printhead*. This part of a printer contains many small nozzles (usually 100 to 200) that spray the ink in small droplets onto the page. Many times, the printhead is part of the *ink cartridge*, which contains a reservoir of ink and the printhead in a removable package. Most color inkjet printers include multiple printheads. Either there will be one for the black cartridge and one for the color one, or there will be one for each of the *CMYK* (cyan, magenta, yellow, and black) print inks. The print cartridge must be replaced as the ink supply runs out.

Inside the ink cartridge are several small chambers. At the top of each chamber are a metal plate and a tube leading to the ink supply. At the bottom of each chamber is a small pinhole. These pinholes are used to spray ink on the page to form characters and images as patterns of dots, similar to the way a dot-matrix printer works but with much higher resolution.

There are two methods of spraying the ink out of the cartridge. Hewlett-Packard (HP) popularized the first method: when a particular chamber needs to spray ink, an electric signal is sent to the heating element, energizing it. The elements heat up quickly, causing the ink to vaporize. Because of the expanding ink vapor, the ink is pushed out of the pinhole and forms a bubble. As the vapor expands, the bubble eventually gets large enough to break off into a droplet. The rest of the ink is pulled back into the chamber by the surface tension of the ink. When another drop needs to be sprayed, the process begins again. The second method, developed by Epson, uses a piezoelectric element (either a small rod or a unit that looks like a miniature drumhead) that flexes when energized. The outward flex pushes the ink from the nozzle; on the return, it sucks more ink from the reservoir.

When the printer is done printing, the printhead moves back to its maintenance station. The *maintenance station* contains a small suction pump and ink-absorbing pad. To keep the ink flowing freely, before each print cycle the maintenance station pulls ink through the ink nozzles using vacuum suction. The pad absorbs this expelled ink. The station serves two functions: to provide a place for the printhead to rest when the printer isn't printing and to keep the printhead in working order.

Head Carriage, Belt, and Stepper Motor

Another major component of the inkjet printer is the head carriage and the associated parts that make it move. The *printhead carriage* is the component of an inkjet printer that moves back and forth during printing. It contains the physical as well as electronic connections for the printhead and (in some cases) the ink reservoir. Figure 5.3 shows an example of a head carriage. Note the clips that keep the ink cartridge in place and the electronic connections for the ink cartridge. These connections cause the nozzles to fire, and if they aren't kept clean, you may have printing problems.

FIGURE 5.3 A printhead carriage (holding two ink cartridges) in an inkjet printer

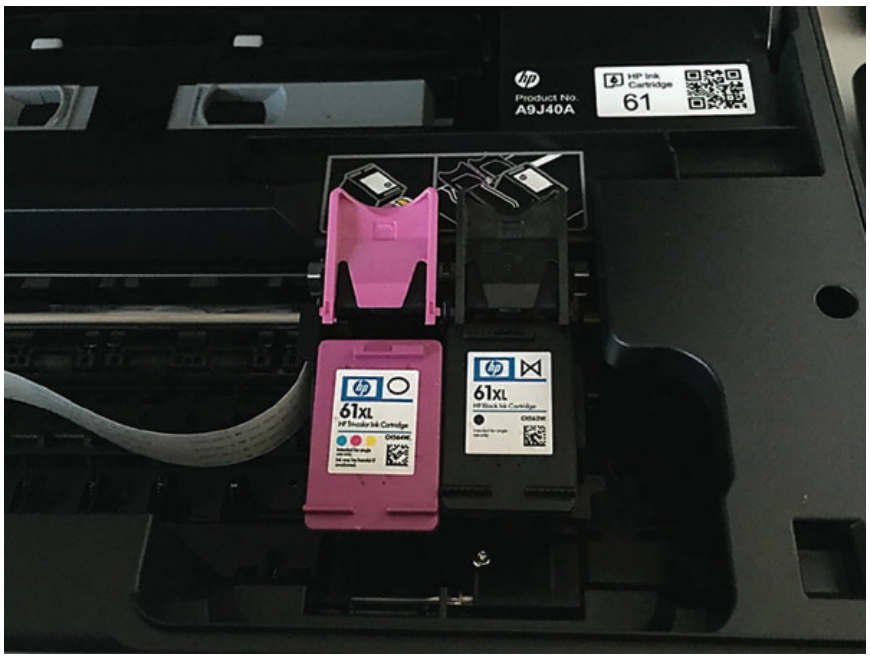

The stepper motor and belt make the printhead carriage move. A *stepper motor* is a precisely made electric motor that can move in the same very small increments each time it is activated. That way, it can move to the same position(s) time after time. The motor that makes the printhead carriage move is also often called the *carriage motor* or *carriage stepper motor*. Figure 5.4 shows an example of a stepper motor.

FIGURE 5.4 A carriage stepper motor

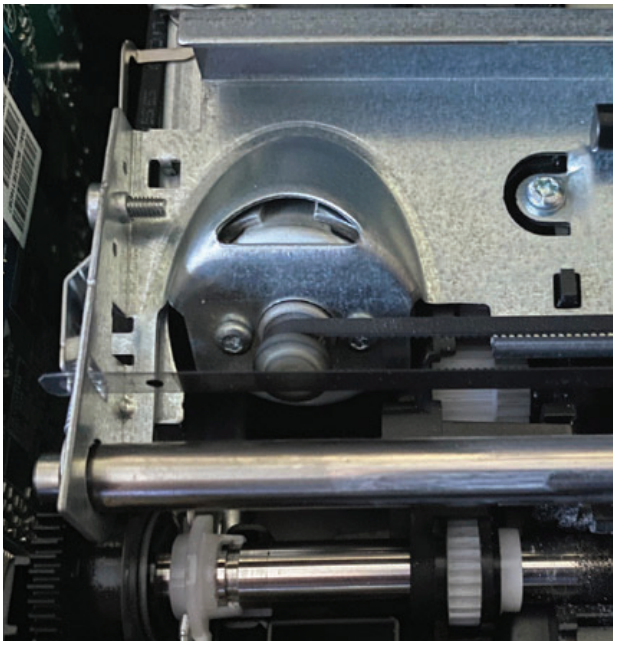

In addition to the motor, a belt is placed around two small wheels or pulleys and attached to the printhead carriage. This belt, called the *carriage belt*, is driven by the carriage motor and moves the printhead back and forth across the page while it prints. To keep the printhead carriage aligned and stable while it traverses the page, the carriage rests on a small metal *stabilizer bar*. Figure 5.5 shows the entire system—the stepper motor, carriage belt, stabilizer bar, and printhead carriage. We will cover inkjet printer maintenance in the "Performing Appropriate Printer Maintenance" section later in this chapter. For now, though, know that if any part of the print carriage assembly breaks in an inkjet printer, it will likely be faster and less expensive to just go buy a new printer rather than trying to fix it. The only maintenance you should need to do is clean the printheads, replace ink cartridges, and clear paper jams.

Paper Feed Mechanism

In addition to getting the ink onto the paper, the printer must have a way to get the paper into the printer. That's where the *paper feed mechanism* (or *feeder*) comes in. The paper feed mechanism picks up paper from the paper drawer and feeds it into the printer. This component consists of several smaller assemblies. First are the *pickup rollers* (see Figure 5.6), which

are one or more rubber rollers with a slightly grippy texture; they rub against the paper as they rotate and feed the paper into the printer. They work against small cork or rubber patches known as *separation pads* (see Figure 5.7), which help keep the rest of the paper in place so that only one sheet goes into the printer. The pickup rollers are turned on a shaft by the *pickup stepper motor*.

FIGURE 5.5 Carriage stepper motor, carriage belt, stabilizer bar, and printhead carriage in an inkjet printer

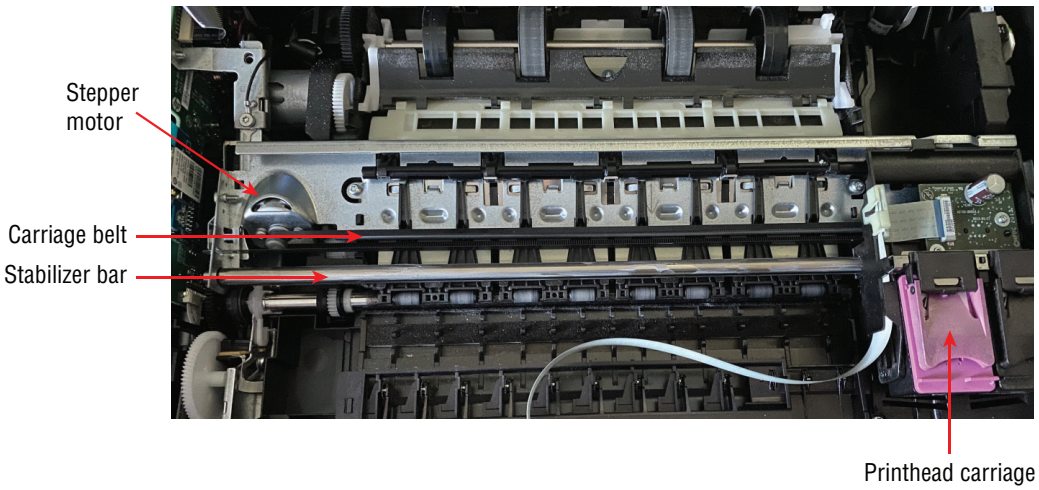

FIGURE 5.6 Inkjet pickup roller (center darker roller)

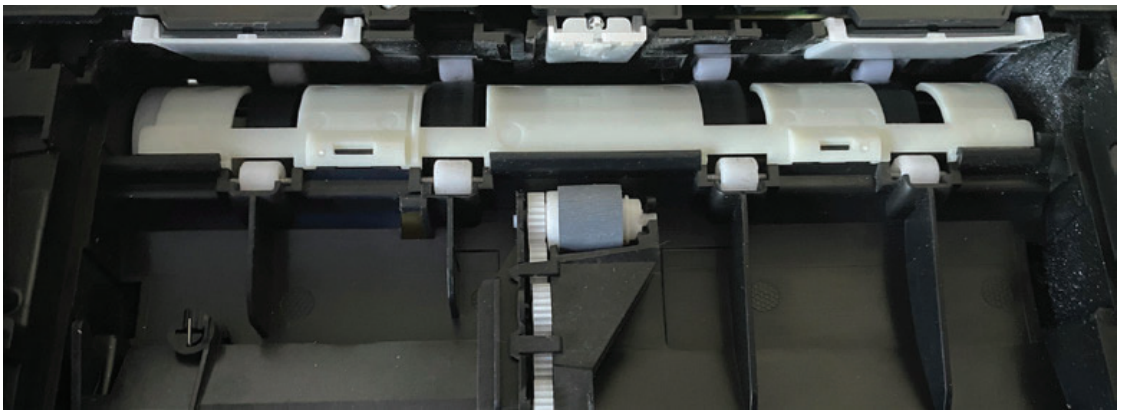

Clean pickup rollers (and other rubber rollers) with mild soap and water, not with alcohol. Alcohol can dry out the rollers, making them brittle and ineffective.

Sometimes the paper that is fed into an inkjet printer is placed into a *paper tray*, which is simply a small plastic tray in the front of the printer that holds the paper until it is fed into the printer by the paper feed mechanism. On smaller printers, the paper is placed vertically into a *paper feeder* at the back of the printer; it uses gravity, in combination with feed rollers and separation pads, to get the paper into the printer. No real rhyme or reason dictates which manufacturers use these different parts; some models use them, and some don't. Generally, more expensive printers use paper trays because they hold more paper. Figure 5.8 shows an example of a paper tray on an inkjet printer.

FIGURE 5.7 Inkjet separation pads

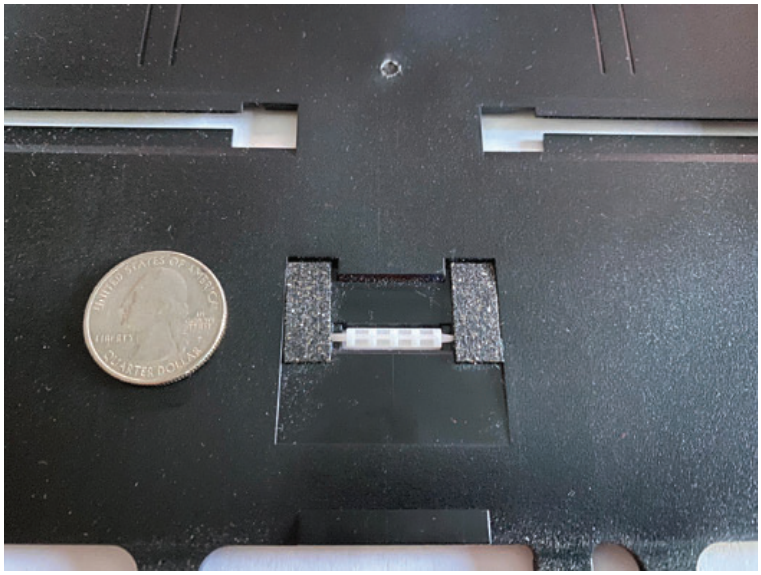

FIGURE 5.8 A paper tray on an inkjet printer

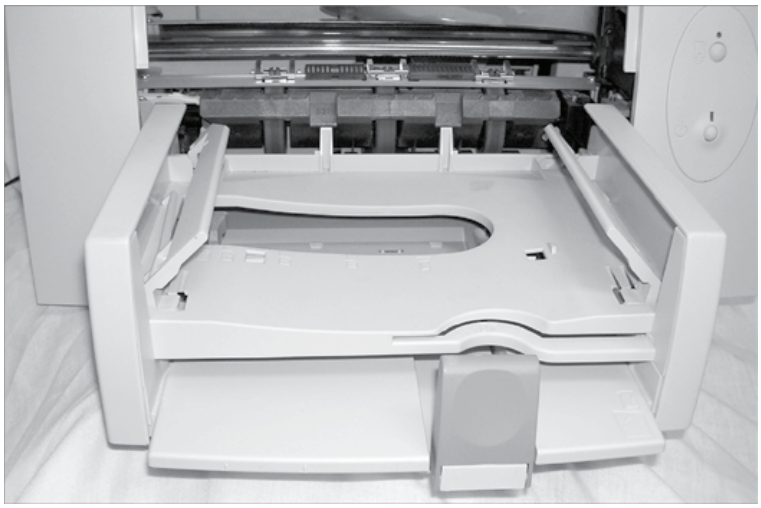

Next are the *paper feed sensors*. These components tell the printer when it is out of paper as well as when a paper jam has occurred during the paper feed process. Figure 5.9 shows an example of a paper feed sensor. Finally, there is the *duplexing assembly*, which allows for printing on both sides of a page. After the first side is printed and has a few seconds to dry, the duplexing assembly will pull the paper back in and flip it over to print the second side. Not all inkjet printers have one, but those that do will usually have them near the back bottom of the printer. Figure 5.10 shows the duplexing assembly rollers. This area is the most likely one to incur a paper jam, so there's usually a removable panel inside the printer for easy access to clear the problem.

FIGURE 5.9 A paper feed sensor on an inkjet printer

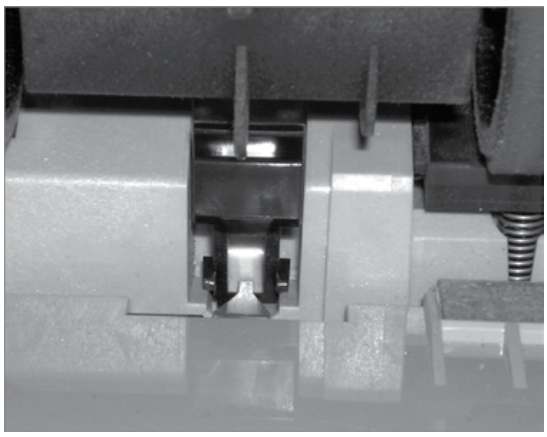

FIGURE 5.10 Duplexing assembly rollers

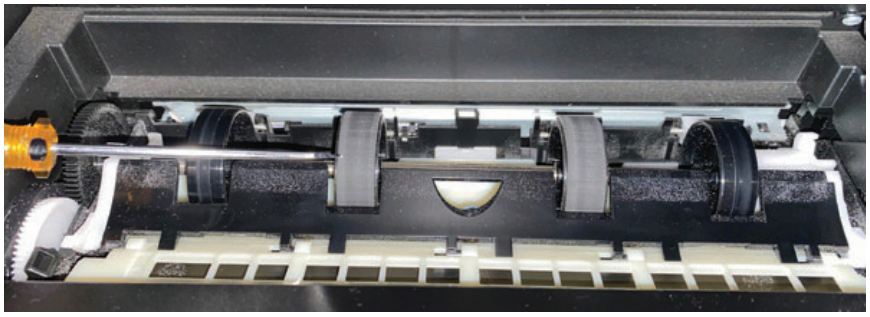

Being able to identify the parts of an inkjet printer is an important skill for an A+ candidate. In Exercise 5.1, you will identify the parts of an inkjet printer. For this exercise, you'll need an inkjet printer.

Many of the exercises in this chapter require printer hardware. If you don't have the proper hardware on which to practice, you can find many helpful videos online showing how to identify parts and install and perform maintenance on printers. Sites such as `IFixit.com` and YouTube `.com` are good places to start.

EXERCISE 5.1

Identifying the Parts of an Inkjet Printer

1. Unplug the inkjet printer from the power source and the computer.
2. Open the top cover to expose the inner print mechanism.
3. Locate and identify the paper tray.
4. Locate and identify the paper feed sensor.
5. Locate and identify the pickup roller(s).
6. Locate and identify the separation pad(s).
7. Locate and identify the printhead, carriage assembly, and belt.
8. Locate and identify the ink cartridge(s).

Control, Interface, and Power Circuitry

The final component group is the electronic circuitry for printer control, printer interfaces, and printer power. The *printer control circuits* are usually on a small circuit board that contains all the circuitry to run the stepper motors the way the printer needs them to work (back and forth, load paper and then stop, and so on). These circuits are also responsible for monitoring the health of the printer and for reporting that information back to the PC.

The second power component, the *interface circuitry* (commonly called a *port*), makes the physical connection to whatever signal is coming from the computer (USB, Ethernet, wireless, etc.) and also connects the physical interface to the control circuitry. The interface circuitry converts the signals from the interface into the data stream that the printer uses.

The last set of circuits the printer uses is the *power circuits*. Essentially, these conductive pathways convert 110V (in the United States) or 220V (in most of the rest of the world) from a standard wall outlet into the voltages that the inkjet printer uses, usually 12V and 5V, and distribute those voltages to the other printer circuits and devices that need it. This is accomplished through the use of a *transformer*. A transformer, in this case, takes the 110V AC current and changes it to 12V DC (among others). This transformer can be either internal (incorporated into the body of the printer) or external. Either design can be used in today's inkjets, although the integrated design is preferred because it is simpler and doesn't show the bulky transformer.

The Inkjet Printing Process

Before you print to an inkjet printer, you must ensure that the device is calibrated. *Calibration* is the process by which a device is brought within functional specifications. For example, inkjet printers need their printheads aligned so that they print evenly and don't print funny-looking letters and unevenly spaced lines. The process is part of the installation for all inkjet printers. Printers will typically run a calibration routine every time you install new ink cartridges. You will only need to manually initiate a calibration if the printing alignment appears off.

Just as with other types of printing, the inkjet printing process consists of a set of steps that the printer must follow in order to put the data onto the page being printed. The following steps take place whenever you click the Print button in your favorite software (like Microsoft Word or Google Chrome):

1. You click the Print button (or similar) that initiates the printing process.
2. The software from which you are printing sends the data to be printed to the printer driver that you have selected.

The function and use of the printer driver and page-description language are discussed later in this chapter.

3. The printer driver uses a page-description language to convert the data being printed into the format that the printer can understand. The driver also ensures that the printer is ready to print.
4. The printer driver sends the information to the printer via whatever connection method is being used (USB, Ethernet, wireless, and so on).
5. The printer stores the received data in its onboard print buffer memory.

   A *print buffer* is a small amount of memory (typically 512 KB to 16 MB) used to store print jobs as they are received from the printing computer. This buffer allows several jobs to be printed at once and helps printing to be completed quickly.
6. If the printer has not printed in a while, the printer's control circuits activate a cleaning cycle.

   A *cleaning cycle* is a set of steps the inkjet printer goes through to purge the printheads of any dried ink. It uses a special suction cup and sucking action to pull ink through the printhead, dislodging any dried ink or clearing stuck passageways.
7. Once the printer is ready to print, the control circuitry activates the paper feed motor.

   This causes a sheet of paper to be fed into the printer until the paper activates the paper feed sensor, which stops the feed until the printhead is in the right position and the leading edge of the paper is under the printhead. If the paper doesn't reach the paper feed sensor in a specified amount of time after the stepper motor has been activated, the Out Of Paper light is turned on and a message is sent to the computer.
8. Once the paper is positioned properly, the printhead stepper motor uses the printhead belt and carriage to move the printhead across the page, little by little.
9. The motor is moved one small step, and the printhead sprays the dots of ink on the paper in the pattern dictated by the control circuitry.

   Typically, this is either a pattern of black dots or a pattern of CMYK inks that are mixed to make colors.

10. Then the stepper motor moves the printhead another small step; the process repeats all the way across the page.

    This process is so quick, however, that the entire series of starts and stops across the page looks like one smooth motion.

11. At the end of a pass across the page, the paper feed stepper motor advances the page a small amount. Then the printhead repeats step 8.
12. Depending on the model, either the printhead returns to the beginning of the line and prints again in the same direction or it moves backward across the page so that printing occurs in both directions. This process continues until the page is finished.
13. Once the page is finished, the feed stepper motor is actuated and ejects the page from the printer into the output tray. (On printers with a duplexing assembly, the paper is only partially ejected so the duplexing assembly can grab it, pull it back in, and flip it over for printing on the second side.)

    If more pages need to print, the process for printing the next page begins again at step 7.

14. Once printing is complete and the final page is ejected from the printer, the printhead is *parked* (locked into rest position) and the print process is finished.

Laser Printers

Laser printers and inkjet printers are referred to as *page printers* because they receive their print job instructions one page at a time rather than receiving instructions one line at a time. There are two major types of page printers that use the electrophotographic (EP) imaging process. The first uses a laser to scan the image onto a photosensitive drum, and the second uses an array of light-emitting diodes (LEDs) to create the image on the drum. Even though they write the image in different ways, both types still follow the laser printer imaging process.

Xerox, Hewlett-Packard, and Canon were pioneers in developing the laser printer technology we use today. Scientists at Xerox developed the electrophotographic (EP) imaging process in 1971. HP introduced the first successful desktop laser printer in 1984, using Canon hardware that used the EP process. This technology uses a combination of static electric charges, laser light, and a black powdery ink-like substance called *toner*. Printers that use this technology are called EP process laser printers, or just *laser printers*. Every laser printer technology has its foundations in the EP printer imaging process.

Let's discuss the basic components of the EP laser printer and how they operate so that you can understand the way an EP laser printer works.

In past versions of the A+ exam, laser printer components were explicitly listed in the exam objectives. With Exam 220-1201 they have been removed. Much like we did for inkjet printers, we've chosen to cover them here anyway because of two reasons. First, Exam Objective 5.6, "Given a scenario, troubleshoot printer issues," still wants you to troubleshoot

laser printer problems, and you may need to troubleshoot them in real life. It's challenging to do that if you don't know how the laser printer is supposed to work in the first place. (As noted earlier, troubleshooting is covered in Volume 1 Chapter 13.) Second, if a question that requires knowledge of laser printer parts slips through to this version of the exam, you will be ready!

Basic Components

Most printers that use the EP imaging process contain nine standard assemblies: the toner cartridge, laser scanner, high-voltage power supply, DC power supply, paper transport assembly (including paper-pickup rollers and paper-registration rollers), transfer corona, fusing assembly, printer controller circuitry, and ozone filter. Let's discuss each of the components individually, along with a duplexing assembly, before we examine how they all work together to make the printer function.

The Toner Cartridge

The EP toner cartridge (see Figure 5.11), as its name suggests, holds the toner. Toner is a black carbon substance mixed with polyester resins to make it flow better and iron oxide particles to make it sensitive to electrical charges. These two components make the toner capable of being attracted to the photosensitive drum and of melting into the paper. In addition to these components, toner contains a medium called the *developer* (also called the *carrier*), which carries the toner until it is used by the EP process.

FIGURE 5.11 An EP toner cartridge

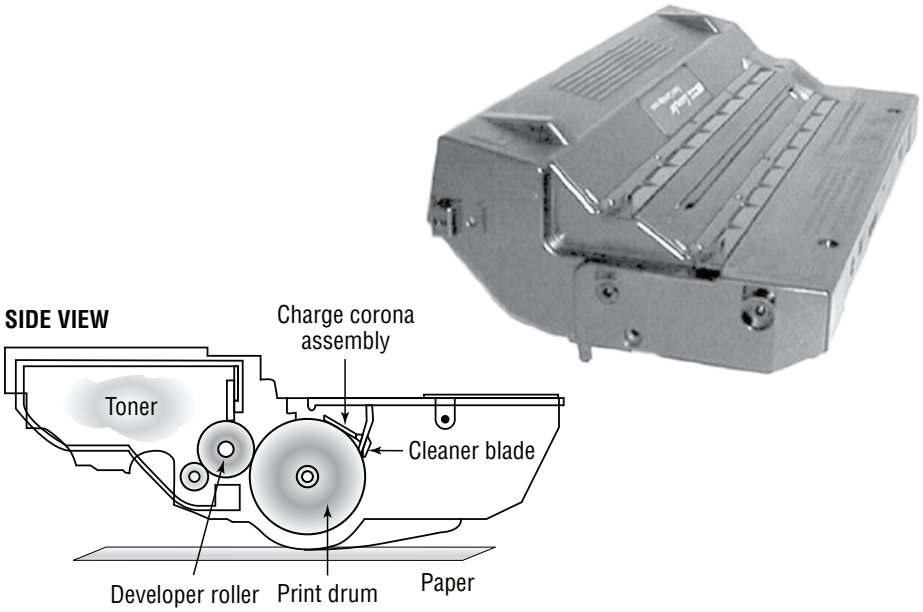

The toner cartridge also contains the EP print drum. This drum is coated with a photosensitive material that can hold a static charge when not exposed to light but *cannot* hold a charge when it *is* exposed to light—a curious phenomenon and one that EP printers exploit for the purpose of making images. Finally, the drum assembly contains a cleaning blade that continuously scrapes the used toner off the photosensitive drum to keep it clean.

Exposing a photosensitive drum to dust or light can damage it, but touching it will most likely render the drum inoperable! It's best to just not mess around with them.

In most laser printers, *toner cartridge* means an EP toner cartridge that contains toner and a photosensitive drum in one plastic case. In some laser printers, however, the toner and photosensitive drum can be replaced separately instead of as a single unit. If you ask for a toner cartridge for one of these printers, all you will receive is a cylinder full of toner. Consult the printer's manual to find out which kind of toner cartridge your laser printer uses.

Never ship a printer anywhere with a toner cartridge installed! The jostling that happens during shipping could cause toner to spill out of the cartridge and all over the inside of the printer. This will be a huge mess to clean up! If the printer is a laser printer, remove the toner cartridge first. You can put it in a sealed, airtight bag to ship, if needed. If it's an LED page printer, there is a method to remove the photosensitive drum and toner cartridge. Also be careful to avoid too much jostling or turning the printer on its side or upside down when unboxing it. (And remember that A+ Exam Objective 3.7 calls out properly unboxing a printer!) Check your manual for details.

The Laser Scanning Assembly

As we mentioned earlier, the EP photosensitive drum can hold a charge if it's not exposed to light. It is dark inside an EP printer, except when the laser scanning assembly shines on particular areas of the photosensitive drum. When it does that, the drum discharges, but only in the area that has been exposed. As the drum rotates, the laser scanning assembly scans the laser across the photosensitive drum, exposing the image onto it. Figure 5.12 shows the laser scanning assembly.

Laser light is damaging to human eyes. Therefore, the laser is kept in an enclosure and will operate only when the laser printer's cover is closed.

FIGURE 5.12 The EP laser scanning assembly (side view and simplified top view)

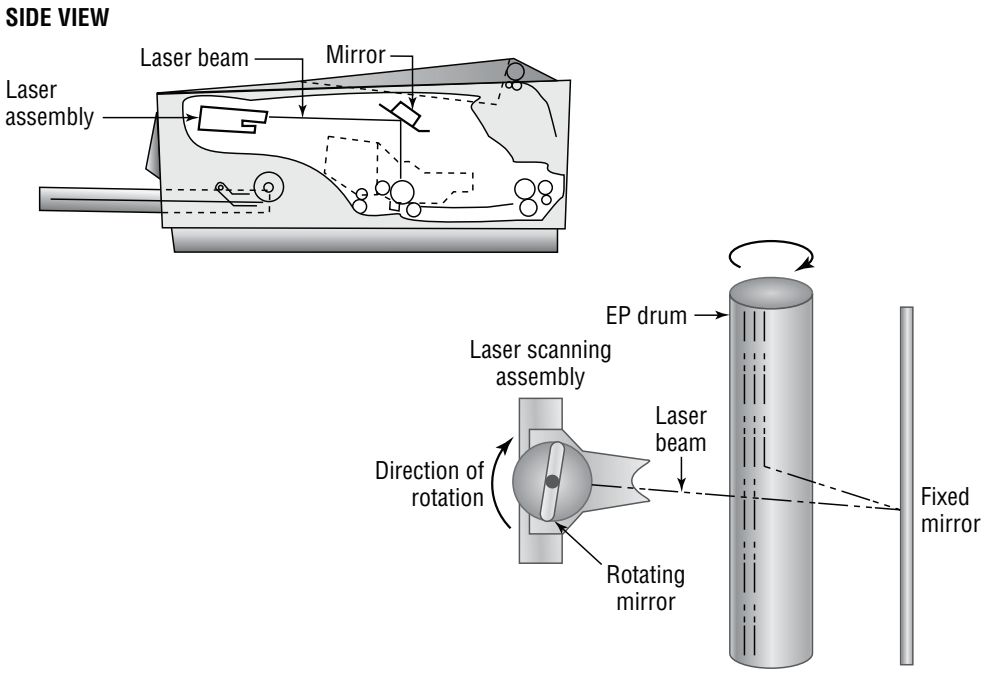

High-Voltage Power Supply

The EP process requires high-voltage electricity. The high-voltage power supply (HVPS) provides the high voltages used during the EP process. This component converts AC current from a standard wall outlet (120V and 60 Hz) into higher voltages that the printer can use. This high voltage is used to energize both the charging corona and the transfer corona.

Anything with the words *high voltage* in it should make you pause before opening a device and getting your hands into it. The HVPS can hurt or kill you if you're working inside a laser printer and don't know what you're doing.

DC Power Supply

The high voltages used in the EP process can't power the other components in the printer (the logic circuitry and motors). These components require low voltages, between +5VDC and +24VDC. The DC power supply (DCPS) converts house current into three voltages: +5VDC and –5VDC for the logic circuitry and +24VDC for the paper transport motors. This component also runs the fan that cools the internal components of the printer.

Paper Transport Assembly

The paper transport assembly is responsible for moving the paper through the printer. It consists of a motor and several rubberized rollers that each performs a different function.

The first type of roller found in most laser printers is the *feed roller*, or *paper pickup roller* (see Figure 5.13). This D-shaped roller, when activated, rotates against the paper and pushes one sheet into the printer. This roller works in conjunction with a special rubber separation pad to prevent more than one sheet from being fed into the printer at a time.

FIGURE 5.13 Paper transport rollers

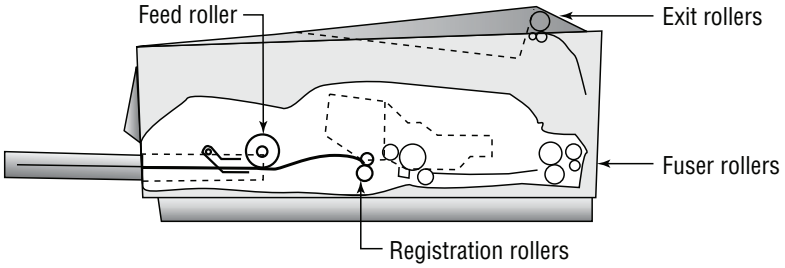

Another type of roller that is used in the printer is the *registration roller* (also shown in Figure 5.13). There are actually two registration rollers, which work together. These rollers synchronize the paper movement with the image-formation process in the EP cartridge. The rollers don't feed the paper past the EP cartridge until the cartridge is ready for it.

Both of these rollers are operated with a special electric motor known as an *electronic stepper motor*. This type of motor can accurately move in very small increments. It powers all the paper transport rollers as well as the fuser rollers.

The Transfer Corona Assembly

When the laser writes (exposes) the images on the photosensitive drum, the toner then sticks to the exposed areas. (We'll cover this later in the "Electrophotographic Imaging Process" section.) How does the toner get from the photosensitive drum onto the paper? The *transfer corona assembly* (see Figure 5.14) is given a high-voltage charge, which is transferred to the paper, which, in turn, pulls the toner from the photosensitive drum.

FIGURE 5.14 The transfer corona assembly

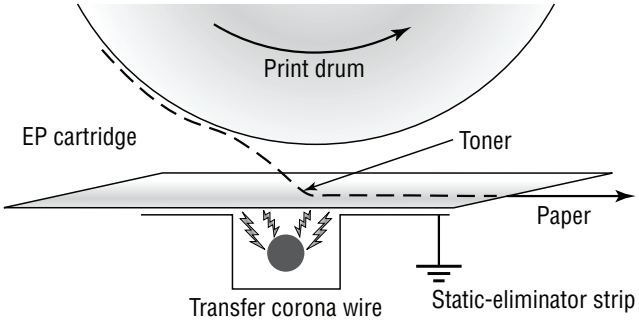

Included in the transfer corona assembly is a *static-charge eliminator strip* that drains away the charge imparted to the paper by the corona. If you didn't drain away the charge, the paper would stick to the EP cartridge and jam the printer.

There are two types of transfer corona assemblies: those that contain a transfer *corona wire* and those that contain a transfer *corona roller*. The transfer corona wire is a small-diameter wire that is charged by the HVPS. The wire is located in a special notch in the floor of the laser printer (under the EP print cartridge). The transfer corona roller performs the same function as the transfer corona wire, but it's a roller rather than a wire. Because the transfer corona roller is directly in contact with the paper, it supports higher speeds. For this reason, the transfer corona wire is used infrequently in laser printers today.

Fusing Assembly

The toner in the EP toner cartridge will stick to just about anything, including paper. This is true because the toner has a negative static charge and most objects have a net positive charge. However, these toner particles can be removed by brushing any object across the page. This could be a problem if you want the images and letters to stay on the paper permanently.

To solve this problem, EP laser printers incorporate a device known as a *fuser* (see Figure 5.15), which uses two rollers that apply pressure and heat to fuse the plastic toner particles to the paper. You may have noticed that pages from either a laser printer or a copier (which uses a similar device) come out warm. This is because of the fuser.

FIGURE 5.15 The fuser

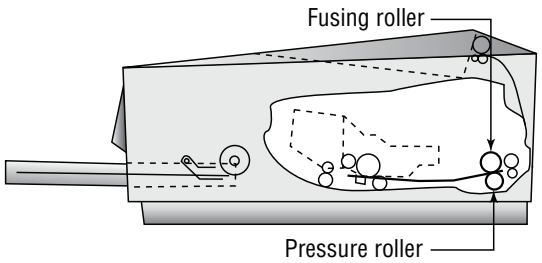

The fuser is made up of three main parts: a halogen heating lamp, a Teflon-coated aluminum-fusing roller, and a rubberized pressure roller. The fuser uses the halogen lamp to heat the fusing roller to between 329° F (165° C) and 392° F (200° C). As the paper passes between the two rollers, the pressure roller pushes the paper against the fusing roller, which melts the toner into the paper.

The fuser can cause severe burns! Be careful when working with it.

Printer Controller Circuitry

Another component in the laser printer that we need to discuss is the *printer controller assembly*. This large circuit board converts signals from the computer into signals for the various assemblies in the laser printer using a process known as *rasterizing*. This circuit board is usually mounted under the printer. The board has connectors for each type of interface and cables to each assembly.

When a computer prints to a laser printer, it sends a signal through a cable to the printer controller assembly. The controller assembly formats the information into a page's worth of line-by-line commands for the laser scanner. The controller sends commands to each of the components, telling them to wake up and begin the EP imaging process.

Ozone Filter

Laser printers use various high-voltage biases inside the case. As anyone who has been outside during a lightning storm can tell you, high voltages create ozone. Ozone is a chemically reactive gas that is created by the high-voltage coronas (charging and transfer) inside the printer. Because ozone is chemically reactive and can severely reduce the life of laser printer components, many older laser printers contain a filter to remove ozone gas from inside the printer as it is produced. This filter must be removed and cleaned with compressed air periodically. (Cleaning it whenever the toner cartridge is replaced is usually sufficient.) Most newer laser printers don't have ozone filters. This is because these printers don't use transfer corona wires but instead use transfer corona rollers, which dramatically reduce ozone emissions.

Duplexing Assembly

Any laser printer worth its money today can print on both sides of the paper (as can some nicer models of inkjet printers, mentioned earlier). This is accomplished through the use of a *duplexing assembly*. Usually located inside or on the back of the printer, the assembly is responsible for taking the paper, turning it over, and feeding back into the printer so the second side can be printed.

Electrophotographic Imaging Process

The *electrophotographic (EP) imaging process* is the process by which an EP laser printer forms images on paper. It consists of seven major steps, each designed for a specific goal. Although many different manufacturers word these steps differently or place them in a different order, the basic process is still the same. Here are the steps in the order in which you should know them:

1. Processing
2. Charging
3. Exposing
4. Developing
5. Transferring
6. Fusing
7. Cleaning

Before any of these steps can begin, however, the controller must sense that the printer is ready to start printing (toner cartridge installed, fuser warmed to temperature, and all covers in place). Printing cannot take place until the printer is in its ready state, usually indicated by an illuminated Ready LED light or a display that says something like `00 READY` (on HP printers). The computer sends the print job to the printer, which begins processing the data as the first step to creating output.

In past versions of the A+ exam, the laser printer imaging process was central to the printing exam objectives. Much like laser printer parts, exam 220-1201 objectives don't list the process. Again, we've chosen to cover them here anyway because of the same two reasons we covered the parts. First, you need to know how to troubleshoot laser printer problems, and second, to prepare you if a question on the process slips through.

Step 1: Processing

The *processing step* consists of two parts: receiving the image and creating the image. The computer sends the print job to the printer, which receives it via its print interface (USB, wireless, etc.). Then, the printer needs to create the print job in such a way that it can accurately produce the output.

If you think back to our discussion of impact printing earlier in this chapter, you might recall that impact printers produce images by creating one strip of dots at a time across the page. Laser printers use the same concept of rendering one horizontal strip at a time to create the image. Each strip across the page is called a *scan line* or a *raster line*.

A component of the laser printer called the *Raster Image Processor* (RIP) manages raster creation. Its responsibility is to generate an image of the final page in memory. How the raster gets created depends on the page-description language that your system is using, such as PostScript (PS) or Printer Control Language (PCL). (We will get into the details of PS and PCL in the "Driver Software" section later in the chapter.) Ultimately, this collection of lines is what gets written to the photosensitive drum and onto the paper.

Step 2: Charging

The next step in the EP process is the *charging step* (see Figure 5.16). In this step, a special wire or roller (called a *charging corona*) within the EP toner cartridge (above the photosensitive drum) gets high voltage from the HVPS. It uses this high voltage to apply a strong, uniform negative charge (around –600VDC) to the surface of the photosensitive drum.

Step 3: Exposing

Next is exposing the drum to the image, often referred to as the *exposing step*. In this step, the laser is turned on and scans the drum from side to side, flashing on and off according to the bits of information that the printer controller sends it as it communicates the individual bits of the image. Wherever the laser beam touches, the photosensitive drum's charge is severely reduced from –600VDC to a slight negative charge (around –100VDC). As the

drum rotates, a pattern of exposed areas is formed, representing the image to be printed. Figure 5.17 shows this process.

FIGURE 5.16 The charging step of the EP process

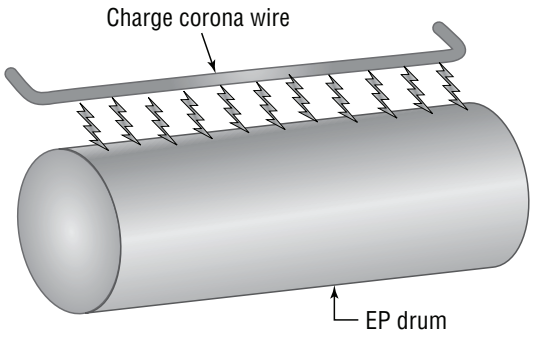

FIGURE 5.17 The exposing step of the EP process

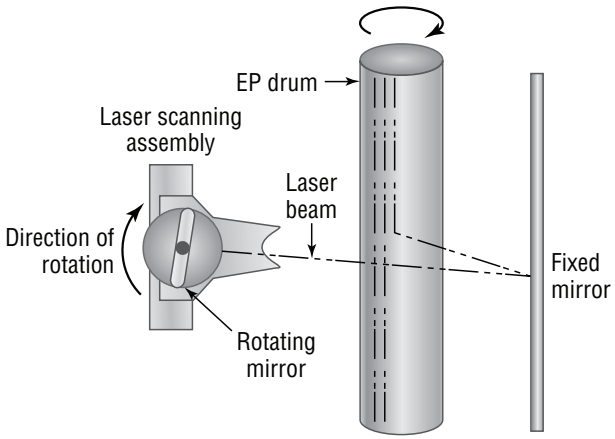

At this point, the controller sends a signal to the pickup roller to feed a piece of paper into the printer, where it stops at the registration rollers.

Step 4: Developing

Now that the surface of the drum holds an electrical representation of the image being printed, its discrete electrical charges need to be converted into something that can be transferred to a piece of paper. The EP process step that accomplishes this is the *developing step* (see Figure 5.18). In this step, toner is transferred to the areas that were exposed in the exposing step.

FIGURE 5.18 The developing step of the EP process

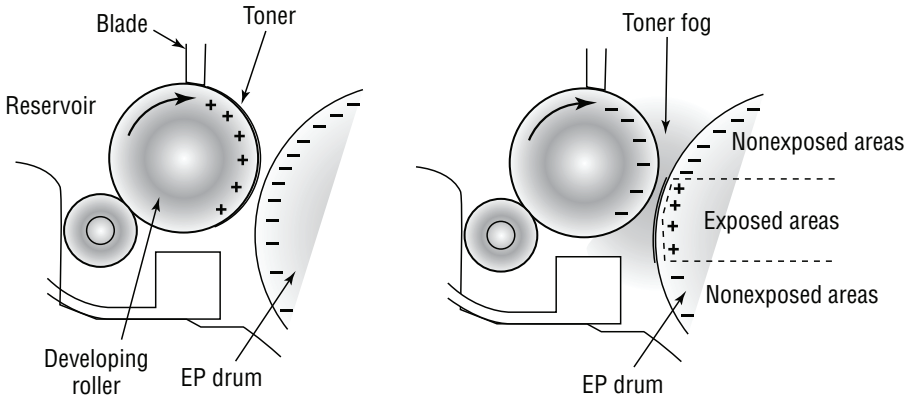

A metallic roller called the *developing roller* inside an EP cartridge acquires a –600VDC charge (called a *bias voltage*) from the HVPS. The toner sticks to this roller because there is a magnet located inside the roller and because of the electrostatic charges between the toner and the developing roller. While the developing roller rotates toward the photosensitive drum, the toner acquires the charge of the roller (–600VDC). When the toner comes between the developing roller and the photosensitive drum, the toner is attracted to the areas that have been exposed by the laser (because these areas have a lesser charge, –100VDC). The toner also is repelled from the unexposed areas (because they are at the same –600VDC charge and like charges repel). This toner transfer creates a fog of toner between the EP drum and the developing roller.

The photosensitive drum now has toner stuck to it where the laser has written. The photosensitive drum continues to rotate until the developed image is ready to be transferred to paper in the next step.

Step 5: Transferring

At this point in the EP process, the developed image is rotating into position. The controller notifies the registration rollers that the paper should be fed through. The registration rollers move the paper underneath the photosensitive drum, and the process of transferring the image can begin; this is the *transferring step*.

The controller sends a signal to the charging corona wire or roller (depending on which one the printer has) and tells it to turn on. The corona wire/roller then acquires a strong *positive* charge (+600VDC) and applies that charge to the paper. Thus charged, the paper pulls the toner from the photosensitive drum at the line of contact between the roller and the paper because the paper and toner have opposite charges. Once the registration rollers move the paper past the corona wire, the static-eliminator strip removes all charge from that line of the paper. Figure 5.19 details this step. If the strip didn't bleed this charge away, the paper would be attracted to the toner cartridge and cause a paper jam.

FIGURE 5.19 The transferring step of the EP process

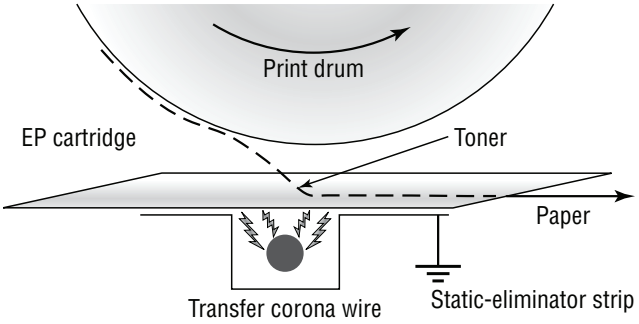

The toner is now held in place by weak electrostatic charges and gravity. It will not stay there, however, unless it is made permanent, which is the reason for the fusing step.

Step 6: Fusing

The penultimate step before the printer produces the finished product is called the *fusing step*. Here the toner image is made permanent. The registration rollers push the paper toward the fuser rollers. Once the fuser grabs the paper, the registration rollers push for only a short time longer. The fuser is now in control of moving the paper.

As the paper passes through the fuser, the 350° F fuser roller melts the polyester resin of the toner, and the rubberized pressure roller presses it permanently into the paper (see Figure 5.20). The paper continues through the fuser and eventually exits the printer.

FIGURE 5.20 The fusing step of the EP process

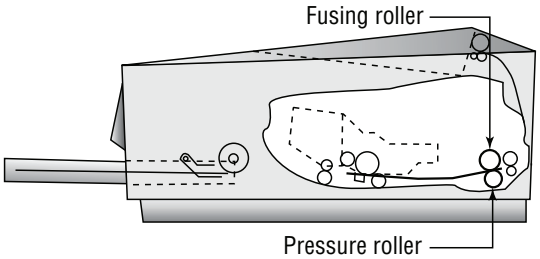

Once the paper completely exits the fuser, it trips a sensor that tells the printer to finish the EP process with the cleaning step.

Step 7: Cleaning

In the last part of the laser imaging process, a rubber blade inside the EP cartridge scrapes any toner left on the drum into a used toner receptacle inside the EP cartridge, and a fluorescent lamp discharges any remaining charge on the photosensitive drum. (Remember that the drum, being photosensitive, loses its charge when exposed to light.) This step is called the *cleaning step* (see Figure 5.21).

FIGURE 5.21 The cleaning step of the EP process

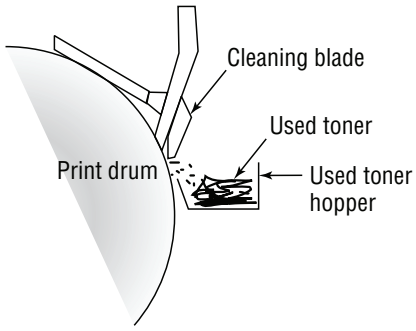

The EP cartridge is constantly cleaning the drum. It may take more than one rotation of the photosensitive drum to make an image on the paper. The cleaning step keeps the drum fresh for each use. If you didn't clean the drum, you would see ghosts of previous pages printed along with your image.

The amount of toner removed in the cleaning process is quite small, and the cartridge will run out of toner before the used toner receptacle fills up. The toner that's in the receptacle is useless because the imaging process has already chemically altered it. In addition, it's considered a hazardous substance. Recycle the print cartridge and don't pay attention to the leftover toner.

At this point, the printer can print another page, and the EP process can begin again.

Summary of the EP Imaging Process

Figure 5.22 provides a diagram of all the parts involved in the EP printing process. Here's a summary of the process, which you should commit to memory:

1. The printer receives and processes the image and stores a page in memory.
2. The printer places a uniform –600VDC charge on the photosensitive drum by means of a charging corona.
3. The laser "paints" an image onto the photosensitive drum, discharging the image areas to a much lower voltage (–100VDC).
4. The developing roller in the toner cartridge has charged (–600VDC) toner stuck to it. As it rolls the toner toward the photosensitive drum, the toner is attracted to (and sticks to) the areas of the photosensitive drum that the laser has discharged.
5. The image is then transferred from the drum to the paper at its line of contact by means of the transfer corona wire (or corona roller) with a +600VDC charge. The static-eliminator strip removes the high positive charge from the paper, and the paper, now holding the image, moves on.

6. The paper then enters the fuser, where a fuser roller and the pressure roller make the image permanent. The paper exits the printer.
7. The printer uses a rubber scraper to clean the photosensitive drum. At that point, it is ready to print the next page or it returns to the ready state.

FIGURE 5.22 The EP imaging process

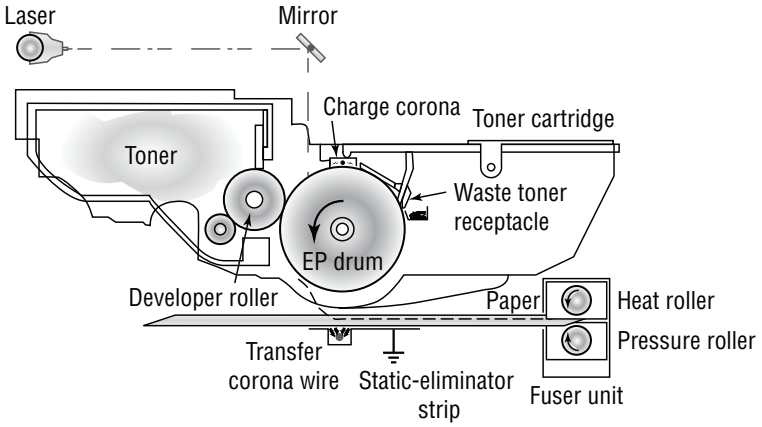

Thermal Printers

The types of printers that you have learned about so far in this chapter account for 90 percent of all paper printers that are used with home or office computers and that you will see as a repair technician. The remaining 10 percent consist of other types of printers that primarily differ by the method they use to put colored material on the paper to represent what is being printed. Examples of these include solid ink, dye sublimation, and thermal printers. Keep in mind that, for the most part, these printers operate like other paper printers in many ways: they all have a paper feed mechanism (sheet-fed or roll); they all require consumables such as ink or toner and paper; they all use the same interfaces, for the most part, as other types of printers; and they are usually about the same size.

Thermal printer concepts you need to know for A+ Exam Objective 3.8 include the feed assembly and special thermal paper. We'll cover the maintenance topics in the "Performing Appropriate Printer Maintenance" section later in this chapter.

Thermal printing technology is primarily used in point-of-sale (POS) terminals. They print on special *thermal paper*, which is a kind of waxy paper that comes on a roll; it's heat-sensitive paper that turns black when heat passes over it. *Thermal printers* work by using a printhead that is the width of the paper. When it needs to print, a heating element heats certain spots on

the printhead. The paper below the heated printhead turns black in those spots. As the paper moves through the printer, the pattern of blackened spots forms an image on the page of what is being printed. Another type of thermal printer uses a heat-sensitive ribbon instead of heat-sensitive paper. A thermal printhead melts wax-based ink from the ribbon onto the paper. These are called *thermal transfer printers* or *thermal wax-transfer printers*.

Thermal direct printers typically have long lives because they have few moving parts. The only unique part that you might not be as familiar with is the paper feed assembly, which often needs to accommodate a roll of paper instead of sheets. The paper is somewhat expensive, doesn't last long (especially if it is left in a very warm place, like a closed car in summer), and produces poorer-quality images than the paper used by most of the other printing technologies.

Installing and Maintaining Printers

Odds are that everyone either owns a printer or has easy access to a printer at a library, work, or some other place. Many retailers and computer manufacturers make it incredibly easy to buy a printer because they often bundle a printer with a computer system as an incentive to get you to buy.

The CompTIA A+ 220-1201 exam will test your knowledge of the procedures to install and maintain printers. In the following sections, we will discuss installing and sharing local and networked printers, implementing network printer security and scan services, and performing printer maintenance.

Many inkjet and laser printers today provide several functions, including copying, scanning, and faxing. These devices go by several names, including multifunctional printer, multifunction device (MFD), multifunction product (MFP), or all-in-one printer.

Installing and Sharing Local Printers

Although every device is different, there are certain accepted methods used for installing almost all of them. The following procedure works for installing many kinds of devices:

1. Choose the proper setup location for the device.
2. Attach the device using a local port (generally USB, but may be something else) and connect the power.
3. Install and update the device driver, update the firmware, and calibrate the device.
4. Configure options and settings.
5. Print a test page.
6. Verify compatibility with the operating system and applications.
7. Educate users about basic functionality.

Before installing any device, read your device's installation instructions. There are exceptions to every rule.

Step 1: Choose the Proper Setup Location for the Device

Before installing the printer, scout out the best location for it. If it's a home-based printer, you may want to choose a convenient but inconspicuous location. In an office setting, having the printer centrally located may save a lot of headaches. Consider the users and how easy or difficult it will be for them to get to the printer. Also consider connectivity. For a wireless printer, how close is it to an access point? (It should be close.) If it's a wired printer, it will need to be near an RJ45 wall jack. Finally, always choose a flat, stable surface. After you determine the location, be sure to properly unbox the device. Most printers will come with a quick setup guide—often a glossy poster-sized printout—that will show you how to properly unbox and connect the device. You may need a box cutter, scissors, or other types of tools to get the box open without destroying it. In general, unboxing steps include:

1. Cutting the tape on the box, taking care as to not damage anything important.
2. Removing any packing materials.
3. Lifting the printer out gently. Some printers are heavy so employ team lifting if needed.
4. Placing it on a stable surface.
5. Removing any protective tape or plastic from the printer itself, including checking inside the paper trays and under the scanner lid (if applicable).
6. Installing paper trays and then ink or toner cartridges according to the manufacturer's instructions.
7. Checking for any additional packaging materials inside the printer that need to be removed.
8. Recycling packing materials and the box.

Always consult the user manual for specific details regarding the printer model.

Exam Objective 3.7, "Given a scenario, deploy and configure multifunction devices/printers and settings," wants you to be able to properly unbox the device and consider setup location.

Step 2: Attach the Device Using a Local Port and Connect the Power

After you have unboxed the printer, with the device powered off, connect it to the host computer. Connecting a printer is generally done with one of three methods: Universal Serial Bus (USB), Ethernet (RJ45), or wireless. Today, the vast majority of local printers are USB, but wireless options supporting Bluetooth or Wi-Fi are common as well.

For Exam Objective 3.7, know the three common connectivity methods: USB, Ethernet, and wireless.

Universal Serial Bus We covered USB in Volume 1 Chapter 3, "Peripherals, Cables, and Connectors," so review there if you need technical details. Given that USB is so popular, it shouldn't surprise you that it's the most common type of wired printer interface as well. Printers often use the USB Type-B connector (it's the square one). The convenience for printers is that it has a higher transfer rate than older serial or parallel connections, and it automatically recognizes new devices. And, of course, USB is physically very easy to connect.

Older printers might connect using a DB-9 serial cable or a DB-25 parallel cable. Most printers with those connections are fortunately retired at this point.

Ethernet Many printers sold today have a wired Ethernet interface that allows them to be hooked directly to an Ethernet cable. These printers have an internal network interface card (NIC) and ROM-based software that allow them to communicate on the network with servers and workstations.

As with any other networking device, the type of network interface used on the printer depends on the type of network to which the printer is being attached. It's likely that the only connection type that you will run into is RJ45 for an Ethernet connection.

Wireless Clearly, people love their Wi-Fi because it enables them to roam around their home or office and still remain connected to one another and to their network. It logically follows that someone came up with the brilliant idea that it would be nice if printers could be that mobile as well—after all, many are on carts with wheels. Most printers today have built-in Wi-Fi interfaces, while others can accept wireless network cards. Wi-Fi–enabled printers support nearly all 802.11 standards (a, b, g, n, ac, ax, be), and the availability of devices will mirror the current popularity of each standard.

The wireless technology that is especially popular among peripheral manufacturers is *Bluetooth*. Bluetooth is a short-range wireless technology; most devices are specified to work within 10 meters (33 feet). Printers such as the HP Sprocket series and OfficeJet 150 mobile printers have Bluetooth capability. We will cover Bluetooth and Wi-Fi in more detail in Volume 1 Chapter 9, "Wireless and SOHO Networking."

When printing with a Bluetooth-enabled device (like a smartphone or tablet) and a Bluetooth-enabled printer, all you need to do is get within range of the device (that is, move closer), select the printer driver from the device, and choose Print. The information is transmitted wirelessly through the air using radio waves and is received by the device.

When Wi-Fi is used to connect printers to a network on a more permanent basis, it is known as *infrastructure mode*. Wi-Fi and Bluetooth can be used to connect a printer temporarily to a single computer (or mobile device), and the connection does not have permanent status. This type of configuration is known as an ad hoc network connection.

Once you have connected the device, connect power to it using whatever supplied power adapter comes with it. Some devices have their own built-in power supply and just need an AC power cord connecting the device to the wall outlet, whereas others rely on an external transformer and power supply. Then, turn on the device.

Step 3: Install and Update the Device Driver, Update the Firmware, and Calibrate the Device

Once you have connected and powered up the device, Windows should automatically recognize it. When it does, a screen will pop up saying that Windows is installing the driver. If a driver is not found, you will be given the option to specify the location of the driver. You can insert the driver media (flash drive, DVD, etc.) that came with the device or point it to the location of the driver files on your hard drive, and the wizard will guide you through the device driver installation.

If Windows fails to recognize the device, you can start the process manually by initiating the Add Printer Wizard to troubleshoot the installation and to install the device drivers. To start the wizard in Windows 11, type **`printer`** into the Windows search bar and then click Printers & Scanners when it appears. Click the Add Device button next to Add A Printer Or Scanner, as shown in Figure 5.23. If a printer is not detected, a link will appear with the words "The printer that I want isn't listed." Click that link, and it will start the wizard shown in Figure 5.24.

FIGURE 5.23 Printers & Scanners

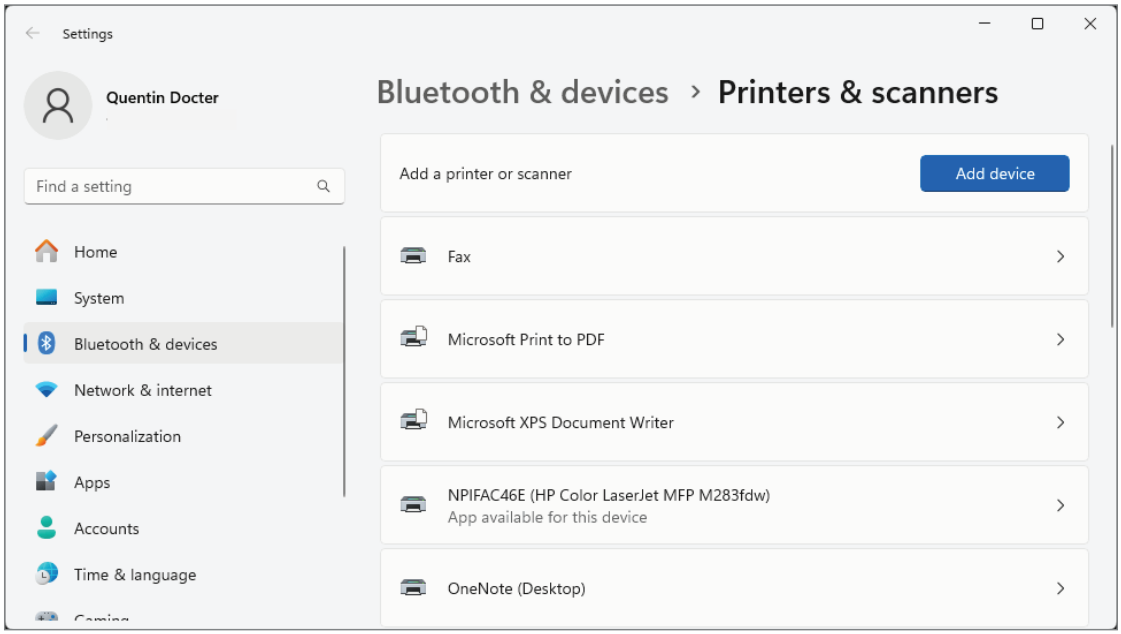

FIGURE 5.24 Add Printer wizard

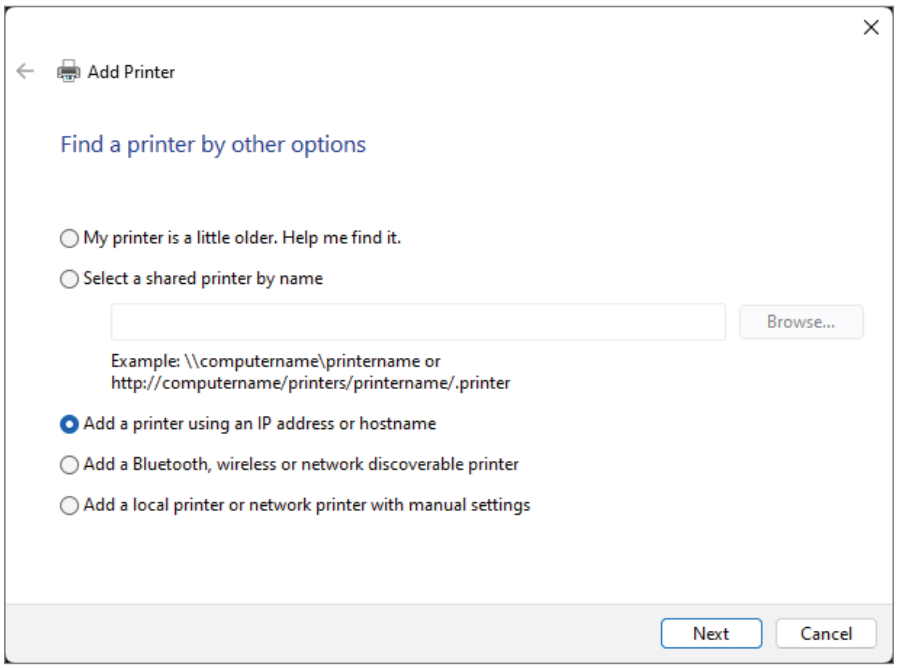
Add Printer
Find a printer by other options
My printer is a little older. Help me find it.
Select a shared printer by name
Browse...
Example: \\computername\printername or
http://computername/printers/printername/.printer
Add a printer using an IP address or hostname
Add a Bluetooth, wireless or network discoverable printer
Add a local printer or network printer with manual settings
Next
Cancel

Driver Software

The *driver* software controls how the printer processes the print job. When you install a printer driver for the printer you are using, it allows the computer to print to that printer correctly. The driver must be written specifically for the operating system the computer is using and for the printer being used. In other words, macOS clients need a different driver than Windows clients need, even to print to the same printer. If you need a refresher on how to install a driver, it's covered in Volume 1 Chapter 2, "Expansion Cards, Storage Devices, and Power Supplies."

If you're working with a Windows-based operating system, Microsoft refers to the software that is installed on the computer and lets you print as the "printer." The physical device where the paper comes out is referred to as the "print device." Here, when we say "printer," we mean the physical device.

When you need to print, you select the printer driver for your printer from a preconfigured list. The driver that you select has been configured for the type, brand, and model of printer, as well as the computer port to which it is connected. You can also select which paper tray the printer should use as well as any other features the printer has (if applicable). Also, each printer driver is configured to use a particular page-description language.

If the wrong printer driver is selected, the computer will send commands in the wrong language. If that occurs, the printer will print several pages full of garbage (even if only one page of information was sent). This "garbage" isn't garbage at all but the printer page-description language commands printed literally as text instead of being interpreted as control commands.

A *page-description language* works just as its name implies: it describes the whole page being printed by sending commands that describe the text as well as the margins and other settings. The controller in the printer interprets these commands and turns them into laser pulses (or pin strikes). Several printer communication languages exist, but the two most common are *PostScript (PS)* and *Printer Control Language (PCL)*.

Hewlett-Packard (HP) developed PCL in 1984, and HP's name for the acronym is Printer Command Language. The CompTIA A+ exam objectives say Printer Control Language. They're not going to ask you to define the acronym on the exam. Be familiar with the basics of what it does, where it's used, and the problems caused by using the wrong page description language.

The first page-description language was PostScript, developed by Adobe in 1984. It made printing graphics fast and simple. Here's how PostScript works. The PostScript printer driver describes the page in terms of "draw" and "position" commands. The page is divided into a very fine grid (as fine as the resolution of the printer). When you want to print a square, a communication like the following takes place:

```
POSITION 1,42%DRAW 10%POSITION 1,64%DRAW10D% . . .
```

These commands tell the printer to draw a line on the page from line 42 to line 64 (vertically). In other words, a page-description language tells the printer to draw a line on the page and gives it the starting and ending points—and that's that. Rather than send the printer the location of each and every dot in the line and an instruction at each and every location to print that location's individual dot, PostScript can get the line drawn with fewer than five instructions. As you can see, PostScript uses commands that are more or less in English. The commands are interpreted by the processor on the printer's controller and converted into the print-control signals.

When HP developed PCL, it was originally intended for use with inkjet printers as a competitor to PostScript. Since then, its role has been expanded to virtually every printer type, and it's a de facto industry standard.

Most newer printers can handle both PS and PCL and will automatically translate for you. Therefore, it's less likely that you'll install the wrong printer driver than it was several years ago.

The main advantage of page-description languages is that they move some of the processing from the computer to the printer. With text-only documents, they offer little benefit. However, with documents that have large amounts of graphics or that use numerous fonts, page-description languages make the processing of those print jobs happen much faster. This makes them an ideal choice for laser printers, although nearly every type of printer uses them.

Exam Objective 3.7 wants you to be familiar with using the appropriate printer drivers, PCL (more common and used for typical office printing) and PostScript (less common and used for professional graphics), and updating device firmware.

Updating Firmware

Once the driver is installed, the device will function. But before anyone uses it, a good next step to take is to check for any firmware upgrades for the printer. Many printers will have a BIOS with firmware, similar to a computer, and manufacturers will occasionally release new versions with new features or security. Firmware updates can be found on the printer manufacturer's website.

Calibrating the Printer

Some print devices must be calibrated before use, particularly inkjets. If the printer requires this step, it will tell you. You'll need to walk through a few steps before the printer will print, but instructions will be provided either on your computer screen or on the printer's display.

When you are working with print media, it is especially important to calibrate all your hardware, including your monitor, scanner, printer, and digital camera, to ensure color matching.

Each manufacturer's process is different, but a typical alignment/calibration works like this:

1. During software installation, the installation wizard asks you if you would like to calibrate now, to which you will respond Yes or OK.
2. The printer prints out a sheet with multiple sets of numbered lines. Each set of lines represents an alignment instance.
3. The software will ask you which set(s) looks the best. Enter the number and click OK or Continue.
4. Some alignment routines end at this point. Others will reprint the alignment page and see if the alignment "took." If not, you can reenter the number of the one that looks the best.
5. Some devices with built-in scanners will ask you to scan the printed image to ensure that calibration looks correct.
6. Click Finish to end the alignment routine.

Real World Scenario

Life Without a Page-Description Language

Page-description languages make printing an efficient process. But what about when they are not used? In situations like this, the computer sends all the instructions the printer needs in a serial stream, like so: Position 1, print nothing; Position 2, strike pins 1 and 3; Position 3, print nothing. This type of description language works well for dot-matrix printers, but it can be inefficient for laser printers. For example, if you wanted to print a page using a standard page-description language and only one character was on the page, there would be a lot of signals wasted on the "print nothing" commands.

With graphics, the commands to draw a shape on the page are relatively complex. For example, to draw a square, the computer (or printer) has to calculate the size of the square and convert that into a lot of "strike pin *x*" (or "turn on laser") and "print nothing" commands. This is where the other types of page-description languages come into the picture.

Step 4: Configure Options and Settings

Once you have installed the drivers and calibrated the device, you can configure any options that you would like for the printer. All the settings and how to change them can be found online or in your user manual.

Where you configure specific printer properties depends a lot on the printer itself. As a rule of thumb, you're looking for the Printer Properties or Printing Preferences applet. In Windows 11, if you open Printers & Scanners, you will see the list of printers that are installed. Clicking an installed printer will reveal several options, shown in Figure 5.25, five of which we'll examine here. There's also a button at the top to remove the device if you choose.

FIGURE 5.25 Printer options

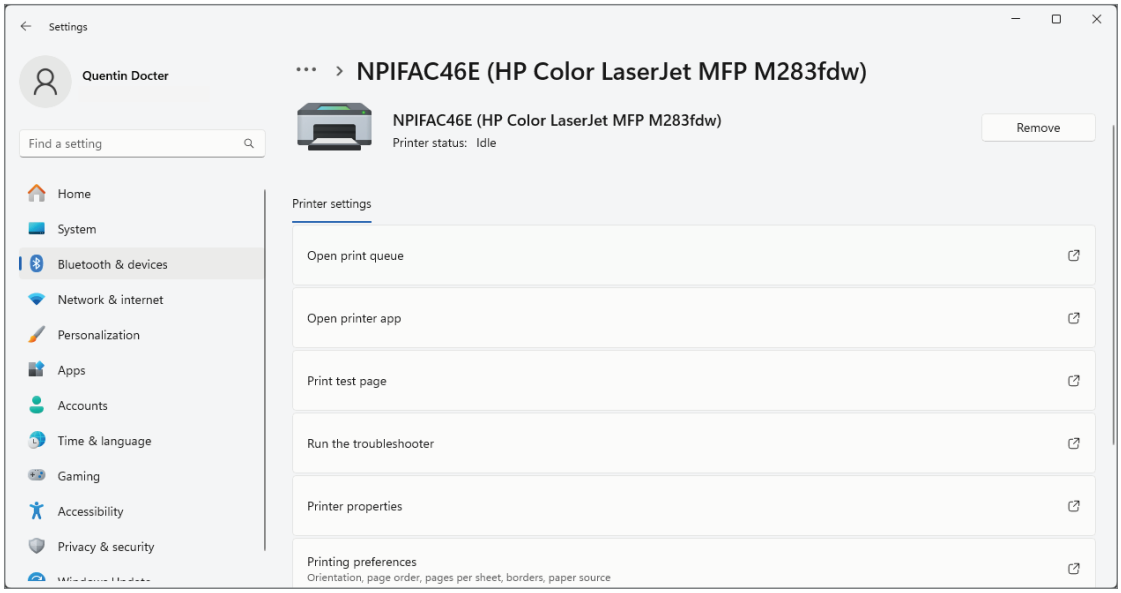

Open Print Queue The print queue (Figure 5.26) is the list of jobs waiting to be printed. This option lets you manage print jobs by pausing, restarting, and canceling print jobs that are in the queue. For example, one could be paused to let others go first, restarted if there was a mishap mid-print, or canceled if no longer needed. Jobs are managed by right-clicking them, or by clicking the three dots in the upper-right corner.

FIGURE 5.26 Print Queue

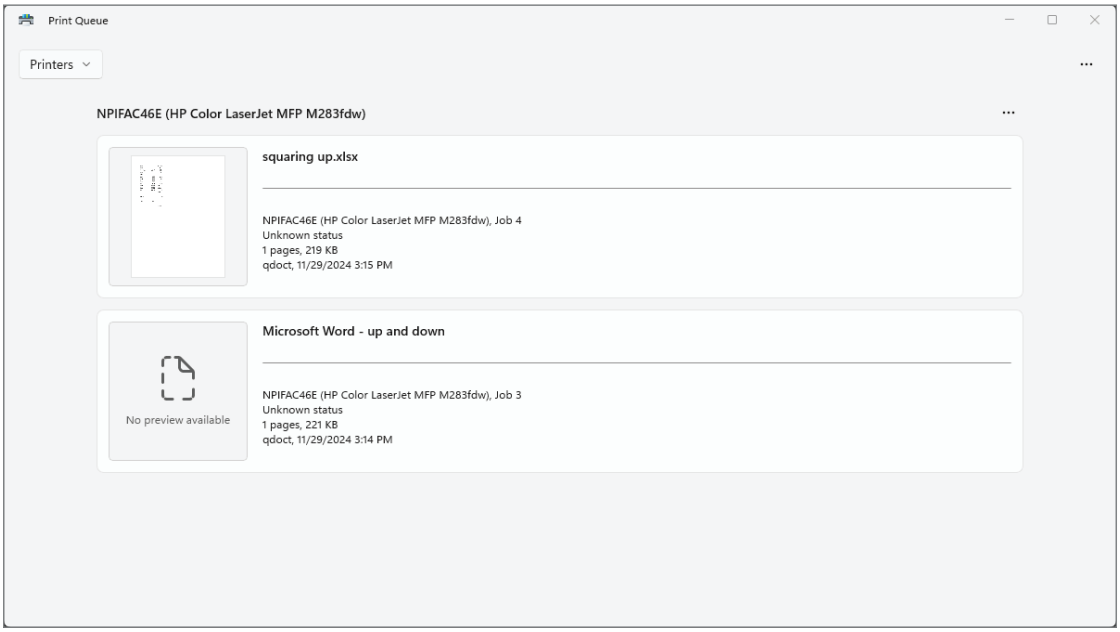
Print Queue
Printers
NPIFAC46E (HP Color LaserJet MFP M283fdw)
squaring up.xlsx
NPIFAC46E (HP Color LaserJet MFP M283fdw), Job 4
Unknown status
1 pages, 219 KB
qdoct, 11/29/2024 3:15 PM
Microsoft Word - up and down
No preview available
NPIFAC46E (HP Color LaserJet MFP M283fdw), Job 3
Unknown status
1 pages, 221 KB
qdoct, 11/29/2024 3:14 PM

Open Printer App The Open printer app will differ depending on the printer, but in general it will give you settings provided by the manufacturer. In the case of HP printers, it will open HP Smart (Figure 5.27), which among other things lets you configure scan options, order ink or toner, or change printer settings.

Print Test Page Printer test page is a good diagnostic tool to check printer connectivity and basic functionality. If a test page won't print, nothing else will either and you should check your connections.

Printer Properties Printer Properties General tab (Figure 5.28) is incredibly useful, and the dialog box will be very similar for any printer that you install. Here the most important tabs are Sharing, Ports, Advanced, and Security. The Sharing tab (Figure 5.29) is where the printer is shared, meaning it's made available to other users on the network. On the Ports tab (Figure 5.30), you can configure the port the printer operates on as well as set up printer pooling. This would be used if you have multiple physical printers that operate under the same printer name.

FIGURE 5.27 HP Smart printer app

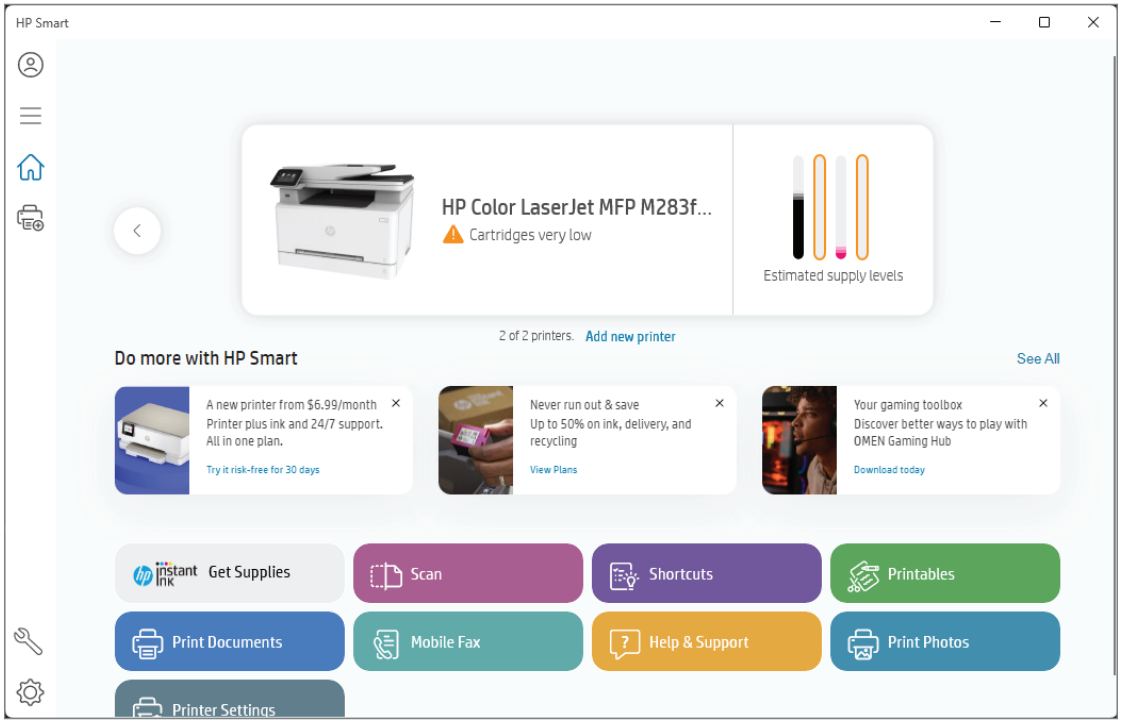

FIGURE 5.28 Printer Properties, General tab

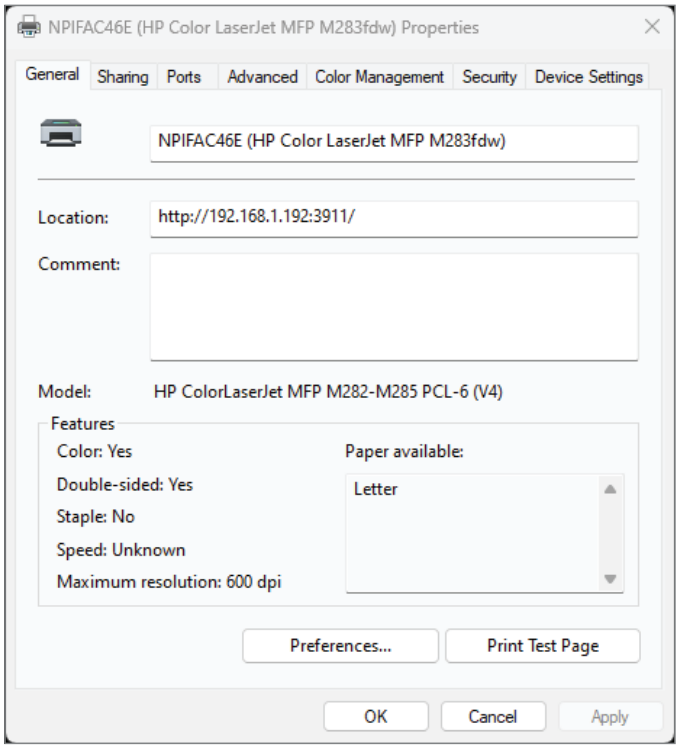

FIGURE 5.29 Sharing tab

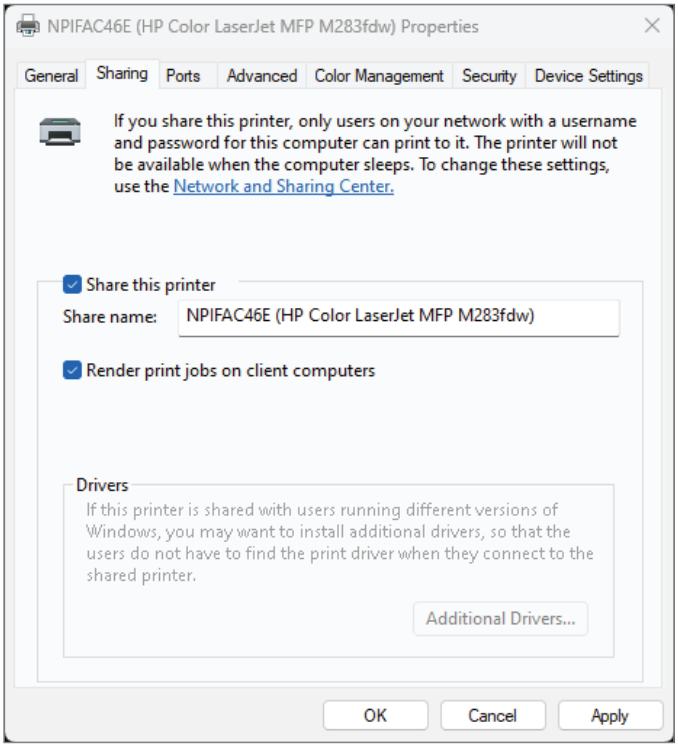
NPIFAC46E (HP Color LaserJet MFP M283fdw) Properties
General
Sharing
Ports
Advanced
Color Management
Security
Device Settings
If you share this printer, only users on your network with a username and password for this computer can print to it. The printer will not be available when the computer sleeps. To change these settings, use the Network and Sharing Center.
Share this printer
Share name:
NPIFAC46E (HP Color LaserJet MFP M283fdw)
Render print jobs on client computers
Drivers
If this printer is shared with users running different versions of Windows, you may want to install additional drivers, so that the users do not have to find the print driver when they connect to the shared printer.
Additional Drivers...
OK
Cancel
Apply

FIGURE 5.30 Ports tab

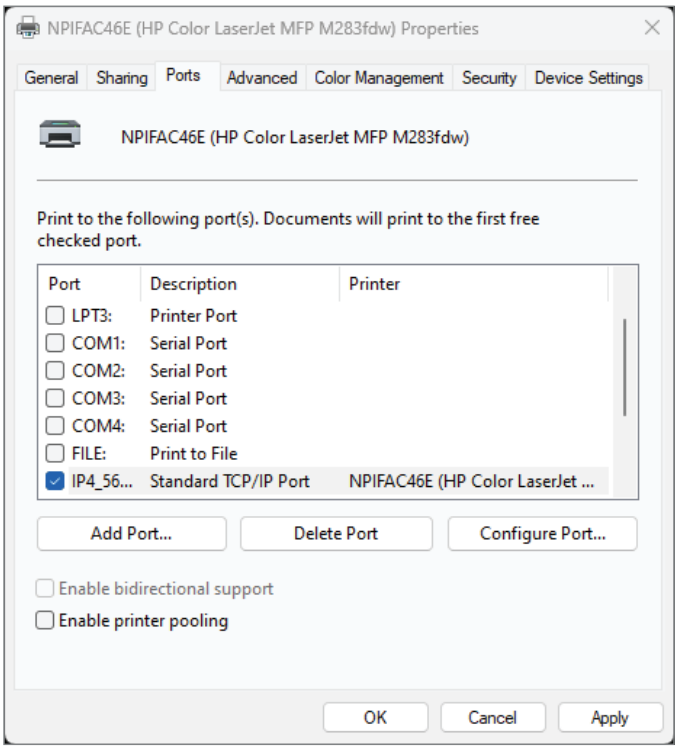

If you're going to configure a printer pool, remember that all of the output can appear on any of the devices that are part of that pool. Make sure that all the printers in that pool are in the same physical location. Otherwise, you will have people wandering all over the office trying to find their printouts. That might be entertaining for you, but not so much for them.

Advanced (Figure 5.31) contains control options. On this tab, you can configure the printer to be available during only certain hours of the day. This might be useful if you're trying to curtail after-hours printing of non–work-related documents, for example. You can also configure the spool settings. For faster printing, you should always spool the jobs instead of printing directly to the printer. However, if the printer is printing garbage, you can try printing directly to it to see if the spooler is causing the problem. (Spooling a job means it gets entirely processed in memory so by the time it starts printing the job goes faster versus needing to pause in the middle to retrieve more of the job.)

FIGURE 5.31 Advanced tab

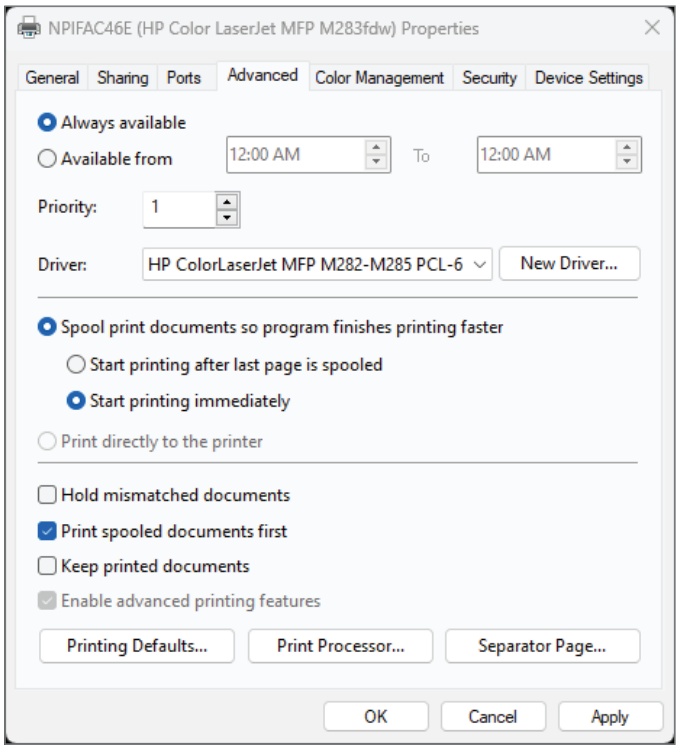

Regarding the options at the bottom, you will always want to print spooled documents first because that speeds up the printing process. If you need to maintain an electronic copy of all printed files, select Keep Printed Documents. Keep in mind that doing so will

eat up a lot of hard disk space and could potentially create a security risk. Finally, there is the Security tab (Figure 5.32), where you control who can print to the printer.

FIGURE 5.32 Security tab

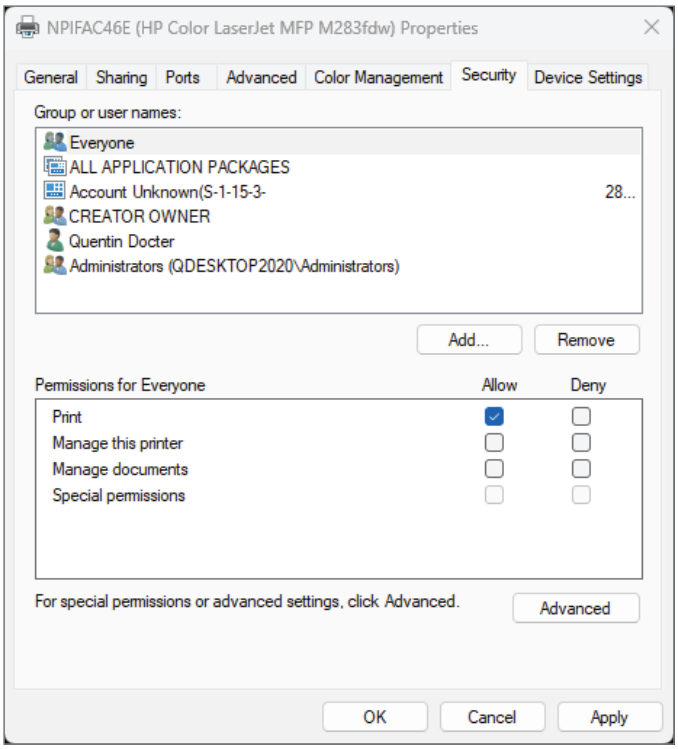

Printing Preferences If you click Printing Preferences in the window shown earlier in Figure 5.25, a window named based on the printer opens. Options in this window (Figure 5.33) will differ based on the installed printer. Here you can configure options such as dual-sided printing and color—of course the available options will depend on what the printer is capable of. Within Printing Preferences you can find configuration settings such as *orientation* (portrait or landscape), *duplex* (double sided), quality, color, and paper tray settings (if applicable) for the printer. Paper tray settings are useful when an MFD supplies multiple sizes of paper, such as letter-sized (8.5″ × 11″) paper and legal-sized (8.5″× 14″) paper in the United States. Or perhaps the MFD has a tray designated for the most common paper size in Europe (and the international standard for office documents), A4, which measures 210mm × 297mm, or 8.3″× 11.7″.

Printer configuration settings, including duplex, orientation, tray settings, and quality, are part of Exam Objective 3.7. Know where to find them!

FIGURE 5.33 Printer-specific options

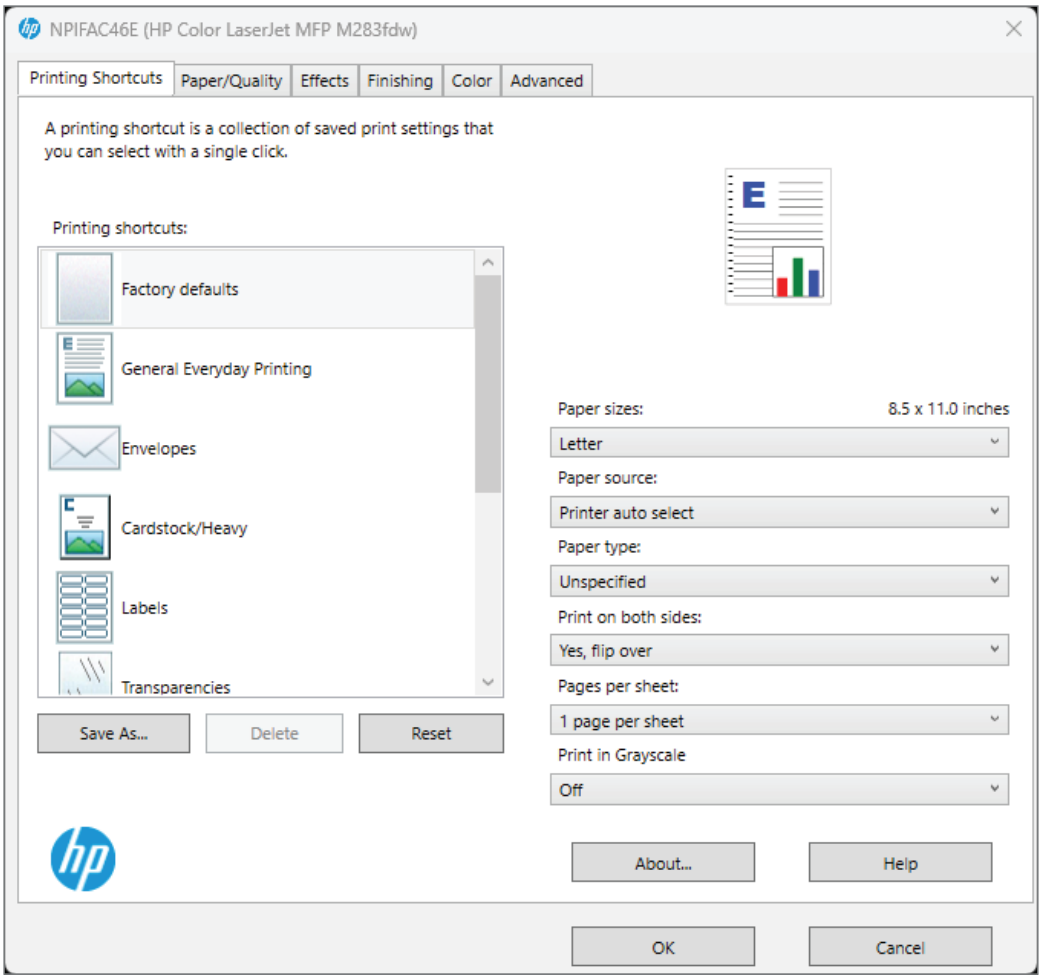

Step 5: Print a Test Page

Once you have configured your printer, you are finished and can print a test page to test its output. Windows has a built-in function for doing just that—you saw links to do so in Figure 5.25 and Figure 5.28. Click the link or button, and Windows will send a test page to the printer. If the page prints, your printer is working. If not, double-check all your connections. If they appear to be in order, then skip ahead to Volume 1 Chapter 13, for troubleshooting tips.

Step 6: Verify Compatibility with Operating System and Applications

Once your printer is installed and you have printed a test page, everything else should work well, right? That's usually true, but it's good practice to verify compatibility with applications before you consider the device fully installed.

With printers, this process is rather straightforward. Open the application you're wondering about and print something. For example, open Microsoft Word, type in some gibberish (or open a real document, if you want), and print it. If you are running non-Microsoft applications (such as a computer-aided drafting program or accounting software) and have questions about their compatibility with the printer, try printing from those programs as well.

Step 7: Educate Users About Basic Functionality

Most users today know how to print, but not everyone knows how to install the right printer or how to print efficiently. This can be a significant issue in work environments.

Say your workplace has 10 different printers, and you just installed number 11. First, your company should use a naming process to identify the printers in a way that makes sense. Calling a printer HPLJ4 on a network does little to help users understand where that printer is in the building. Naming it after its physical location might make more sense.

After installing the printer, offer installation assistance to those who might want to use the device. Show users how to install the printer in Windows (or if printer installation is automated, let them know that they have a new printer and where it is). Also, let users know the various options available on that printer. Can it print double-sided? If so, you can save a lot of paper. Show users how to configure that. Is it a color printer? Do users really need color for rough drafts of documents or presentations? Show users how to print in black and white on a color printer to save the expensive color ink or toner cartridges.

On the printer we've used as an example in this chapter, most of the options involving print output are accessed by clicking Printing Preferences to open the printer-specific settings (shown earlier in Figure 5.33). Duplex (or Print On Both Sides) and Tray settings (or Paper Source) are on the Printing Shortcuts tab. Print Quality (Best, Normal, Draft) is on the Paper/Quality tab, and Orientation (Portrait or Landscape) is set on the Finishing tab. Collation, which enables you to select whether you want it to print pages in order (1, 2, 3. . . 1, 2, 3. . . and so on) or multiple copies of the same page at once (1, 1, 1. . . 2, 2, 2. . . and so forth), is set on the Advanced tab. Sometimes it just takes a bit of poking around to find out where to set the option the user needs.

For the A+ exam, you don't need to memorize which tabs these configuration settings are on. Know what the options do (particularly duplex, orientation, tray settings, and quality), and if you're presented with a screenshot or a series of screenshots, be able to click on the right option to enable the proper feature for the situation.

In Exercise 5.2, we'll step through the process of installing a USB printer in Windows 11.

EXERCISE 5.2

Installing a USB Printer in Windows 11

For this exercise, you will need the following items:

- A USB printer
- A USB printer cable
- The software driver that came with the printer (flash drive, DVD, or downloaded)
- A computer with a free USB port and the ability to access the driver

1. Turn on the computer.
2. Plug the printer into the wall AC outlet and turn it on.
3. Access the driver file to begin the installation process. Typically this is done by inserting the flash drive provided by the manufacturer, and then double-clicking the `setup.exe` or `install.exe` program on the drive. Or, if you downloaded it from the manufacturer's website, find the file and run it from the location where you saved it.
4. Follow the prompts in the installation program to install the driver.
5. Once the software has been installed, plug one end of the USB cable into the printer and the other end into the free USB port. Some installation programs will prompt you for this step.
6. Windows will automatically detect the new printer, install the driver, and configure it.

   Windows will display a balloon in the lower-right corner of the screen that says, "Your hardware is now installed and is ready to use." If Windows doesn't properly detect the printer, open the Add Printer wizard to begin the installation process again, and manually specify the location of the printer driver (such as the flash drive or download location).
7. Print a test page to see if the printer can communicate and print properly.

Real World Scenario

Which Printer Did That Go To?

One of the authors used to work at a satellite office in Salt Lake City for a company whose headquarters were in Houston. Because of printer problems, a new network printer had been installed, which had a different network name from the previous printer.

At the end of the month, one of the accountants printed her monthly reconciliation report, which typically ran about 400 pages. (A hard copy was required for regulatory reasons.)

Puzzled when it didn't come out of the printer, she printed it again. And again. And again. After the fourth failed attempt, and several hours later, she decided to ask someone in IT what the problem was.

It turns out that she had mapped (installed) the new network printer but had gotten a few letters wrong in the printer name. Instead of being at her office, all her print jobs were sent to a printer in the Houston office. And, of course, there were people in Houston trying to print similar reports and who just kept refilling the printer with paper because they didn't want to cut someone else's report off in the middle.

While this wasn't a catastrophic failure, it was annoying. She had unintentionally wasted three reams of paper, the associated toner, and hours of printer life. It wasn't a malicious act, and she was a literate computer user, but it's illustrative of the need to educate and help users with installing and configuring devices. Had the printer been mapped correctly the first time, the waste could have been avoided.

Installing and Sharing Networked Printers

The previous section was about installing a printer attached to your local computer. There are advantages to that approach, such as being able to manage and control your own printer, not to mention having a printer at your own desk. That doesn't happen often in the business world these days!

There are some big disadvantages as well. First, it means that all users who need to print to your device may need local user accounts on your computer, unless you are on a network domain. If so, you will need to manage security for these accounts and the printer. Second, your computer is the print server. The *print server* is the device that hosts the printer and processes the necessary printer commands. This can slow your system down. Third, because your computer is the print server, if for any reason it's turned off, no one will be able to print to that device.

Exam Objective 3.7 lists two subobjectives under public/shared devices. The first is printer share, which simply means a shared printer on the network. The second is print server. Be sure to understand the difference between the two for the exam.

There is another option, though. Instead of needing a specific computer to be the print server, why not make the print server part of the printer itself, or make it a separate network device that hosts the printers? That is exactly the principle behind network printing. Next, we will cover two types of network printing—local network printing and cloud printing—as well as talk about data privacy concerns with printing to public or shared printers.

Local Network Printing

The key to local network printing is that you are moving the print server from your computer to another location, accessible to other users on the network. Therefore, the print server needs a direct attachment to the network, via either a wired (RJ45) or wireless connection. You will find two major varieties of print servers. The first, called an *integrated print server*, is incorporated into the printer itself, and the second is a separate hardware print server. If you are using a stand-alone print server, the printers attach to the print server, either physically or logically. In most cases, if a printer is capable of connecting directly to a network, it has the capability to be its own print server.

Installing and using a networked printer is very similar to installing and using a local printer. You need to ensure that both devices are plugged in, turned on, and attached to the network (either with an RJ45 Ethernet connection or by using wireless). Probably the biggest difference is that when you install it, if it's not detected automatically, you need to tell your computer that you are adding a networked printer instead of a local printer. For example, in Windows 11, when you open the Add Printer utility, you choose from several network printer options, such as Add A Bluetooth, Wireless Or Other Network Discoverable Printer. From there, you will be asked to install the printer driver, just as if the printer were directly attached to your computer. Once it's installed, you use it just as you would use a local printer, including setting the configuration options that we looked at in earlier sections. Every computer on the local network should be able to see and add the printer in the same way.

The print server needs to have drivers available (installed) for all the types of clients that will connect to its printers. For example, if the network has macOS, Linux, and Windows clients, the server will need to have all three drivers. If not, users may not be able to install the printer properly and may not be able to print.

There are a few other ways that you can add shared networked printers: by using TCP, Bonjour, or AirPrint.

TCP Printing

Printers that are network-aware need IP addresses (every device on a network needs a unique address), so it makes sense that you can add a networked printer by using TCP/IP, also known as *TCP printing*. Exercise 5.3 walks you through the general process of installing a TCP printer, using Windows 11 as an example.

EXERCISE 5.3

Installing a TCP/IP Printer in Windows 11

1. Connect the printer to the network and power it on.
2. Configure the printer with an IP address, if it does not already have one.

EXERCISE 5.3 *(continued)*

Most network-aware printers will have their own display panel where you can configure or verify network settings. The IP address needs to be on the same subnet as the computer trying to add the printer.

3. From your Windows 11 computer, start the Add Printer wizard. Open Printers & Scanners, then click Add Device. Even if it finds a printer, ignore that and click the Add Manually link. It will open the Add Printer wizard (shown earlier in Figure 5.24).
4. Select Add A Printer Using A TCP/IP Address Or Hostname, and then click Next.
5. On the next screen, shown in Figure 5.34, you have a few options. One is to let the system automatically detect the device. If so, set the Device Type as Autodetect and click Next. Otherwise, set Device Type to TCP/IP Device and move to the next step. For this exercise, let's do the latter.

FIGURE 5.34 Adding an IP printer

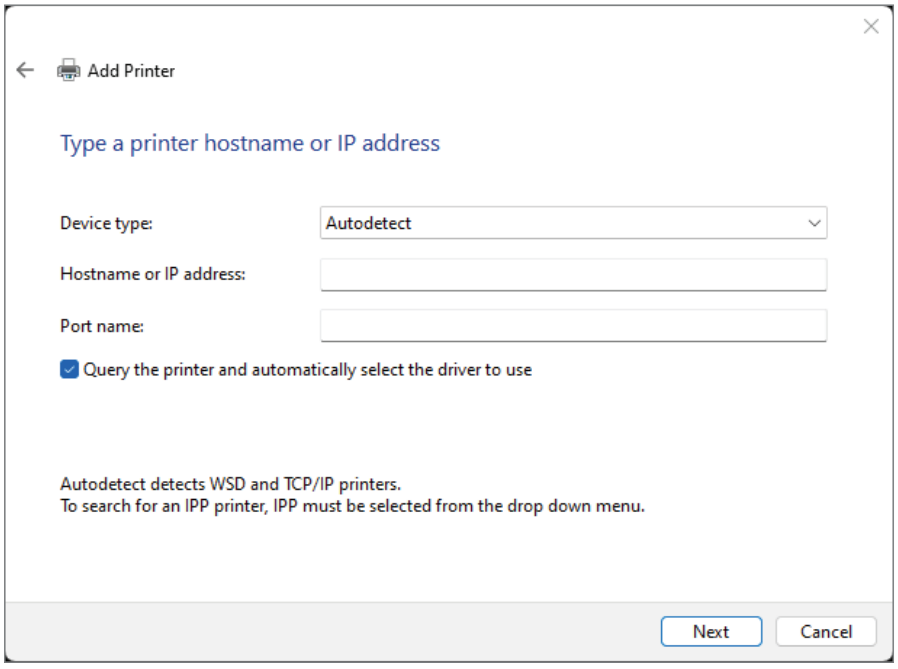
Add Printer
Type a printer hostname or IP address
Device type:
Autodetect
Hostname or IP address:
Port name:
Query the printer and automatically select the driver to use
Autodetect detects WSD and TCP/IP printers.
To search for an IPP printer, IPP must be selected from the drop down menu.
Next
Cancel

6. Enter the IP address (or hostname) of the printer that you want to add, add a port name (it's just for identification), and then click Next.
7. Select the make and model of your printer from the list.
8. You will be prompted to install the driver; continue installing the printer as you would any other local or networked printer.

Some installations will ask you which TCP printing protocol you want to use, RAW or LPR. RAW (also called the Standard TCP/IP Port Monitor) is the default, and it uses TCP port 9100. It also uses the Simple Network Management Protocol (SNMP) for bidirectional communication between the computer and the printer. LPR is older, and the protocol is included for use with legacy systems. It's limited to source ports 721–731 and the destination port 515.

We will discuss ports and port numbers in Volume 1 Chapter 7, "Introduction to TCP/IP."

After the printer is installed, it will appear in your Printers & Scanners window, just as any other printer would.

There are a few advantages to using TCP printing. First, it sends the print jobs directly to the printer, so your system does not need to act as the print server or spend processing time dealing with formatting the print job. Second, it allows clients with different OSs, such as Linux or macOS, to add printers without worrying about intra-OS conflicts.

Bonjour

Apple introduced *Bonjour* in 2002 (then under the name Rendezvous) as an implementation of zero configuration networking. It's designed to enable automatic discovery of devices and services on local networks using TCP/IP as well as to provide hostname resolution. Currently, it comes installed by default on Apple's macOS and iOS operating systems. Bonjour makes it easy to discover and install printers that have been shared by other Bonjour-enabled clients on the network.

Even though Apple developed Bonjour, it does work on other operating systems. For example, it comes with iTunes and the Safari browser, so if you have either of those installed on a Windows computer, odds are that you have Bonjour as well. Once installed, the Bonjour service starts automatically and scans the network looking for shared devices. Exercise 5.4 shows you how to see if Bonjour is installed in Windows.

EXERCISE 5.4

Determining if Bonjour Is Installed in Windows

1. Type **services** into the Windows 11 search bar and then click Services when the system app appears. The Services window will open.
2. Sort the list by name and look for the Bonjour service, as shown in Figure 5.35.

EXERCISE 5.4 *(continued)*

FIGURE 5.35 Bonjour service is running and set to start automatically.

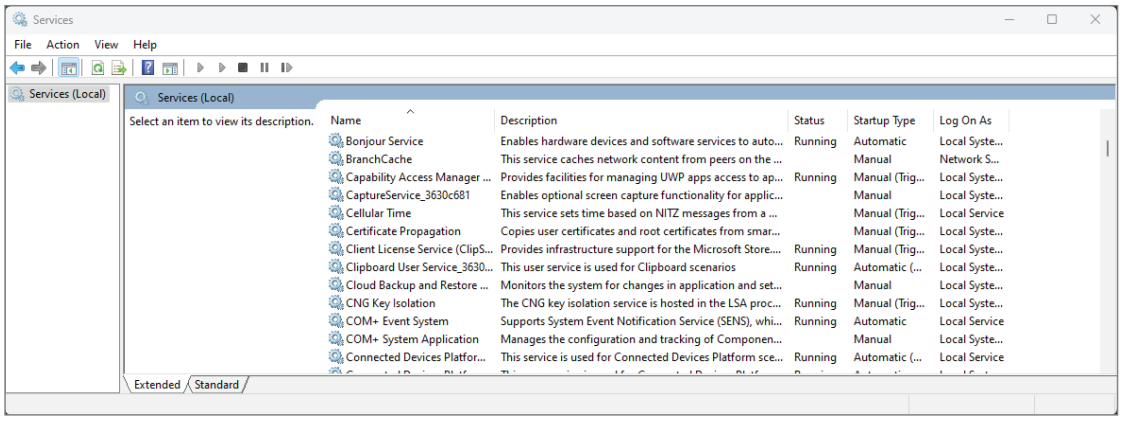

Bonjour works only on a single broadcast domain, meaning that it will not find a printer or other device if it's on the other side of a router from your computer. All major printer manufacturers support Bonjour technology.

If you are using macOS, adding a Bonjour printer is easy. You open System Preferences ➢ Print And Scan, click the plus sign under Printers to open the Add Printer window, and look for the printer on the list. If macOS doesn't have the driver available, you will be asked to provide it. Otherwise, you're done.

In order to add or share a Bonjour printer from Windows, if it doesn't have the service, you need to download Bonjour Print Services for Windows. It's found on Apple's support site at `https://support.apple.com/en-us/106380`.

AirPrint

The one big complaint that Apple aficionados had about Bonjour was that it didn't support printing from iPhones or iPads. In 2010, Apple introduced *AirPrint* to meet that need.

The idea behind AirPrint is quite simple. Mobile devices can automatically detect AirPrint-enabled printers on their local network and print to them without requiring the installation of a driver. To be fair, what Apple really did was eliminate the need for a specific printer driver to be installed on the client and replaced it with the AirPrint concept. Then it was up to the printer manufacturers to develop their own drivers that talked to AirPrint. HP was happy to oblige with its Photosmart Plus series, and other manufacturers soon followed. The list of AirPrint-enabled printers is available at `https://support.apple.com/en-us/HT201311`. From the end-user standpoint, though, no driver is required.

There really is no installation process, and printing is easy. Just be sure that your mobile device is on the same local network as an AirPrint printer. When you attempt to print from your device, select the printer to which you want to print, and it should work.

You can also purchase AirPrint servers, which are small print servers that enable you to connect almost any printer to the network and make them AirPrint-compatible.

Network Printer Security

When printing to a public printer, or one that is shared in a common workspace, there may be data privacy concerns. For example, employees in Human Resources (HR) might need to print confidential personnel files, or someone in the Mergers group might have a highly restricted contract to print. Let's take a look at some security options for networked printers.

User Authentication

Requiring *user authentication* means that users to log in (authenticate) to the printer is one step that can improve printer security. Not all printers have this capability, but most newer laser printers and MFDs designed for office use will. Figure 5.36 shows the front of a Xerox AltaLink MFD. In the center of the picture is a touchscreen display where a user can copy, scan to email, look at printer jobs, configure the device, and use a secure print feature called SafeQ. If a user taps the LogIn button in the upper-left corner, they will be presented with a keyboard so that they can enter their username and password. While functional, this is a bit old-school. An easier way for many users is to scan their work badge (this is called *badging*) on the badge reader at the left of the unit. Doing so will automatically log them in and provide access to secure printing features.

FIGURE 5.36 Xerox AltaLink badge scanner and touchscreen

Exam Objective 3.7 wants you to know four printer security options: user authentication, badging, secured prints, and audit logs.

Secured Prints

Printing a document is usually pretty straightforward. On the computer, the user hits some sort of Print button, chooses the printer to send it to, and presses Print again (or OK, or something similar). Then they get up and walk to the printer to retrieve the hard copy. Most office denizens have been well trained in this process.

Some print jobs might contain sensitive information, though, so the user wants to ensure that they are physically at the printer before it starts printing. Or, perhaps the user has printed to a device in a different building and wants the printer to wait to start printing until they can get there. In cases like these, a *secured prints* feature can be used to hold the print job until the user is ready for it.

Looking back at Figure 5.36, you can see a feature on the bottom center of the touchscreen called SAFEQ Print. YSoft SAFEQ is an industry-standard enterprise print management suite adopted by many organizations, and it works seamlessly with many printers. The administrator sets it up on the printer (and we mean logical printer here, not necessarily the physical device), and when the user prints to that printer, the job waits for them until they authenticate on the physical printer and tell it to start. Figure 5.37 shows the SAFEQ user authentication screen. Once logged in, the user will be presented with the job(s) to print, as shown in Figure 5.38.

FIGURE 5.37 SAFEQ authentication screen

YSoft SafeQ - User authentication
YSoft SafeQ Terminal Embedded
• Use your card or
• enter your username and password
xerox

Audit Logs

Being able to see who used (or perhaps abused) a printer after the fact can come in handy. Some printers have the ability to save a list of documents that have been printed as an *audit log*. The Xerox printer we've used as an example in this section does just that, and the log is shown in Figure 5.39.

FIGURE 5.38 Secured print job

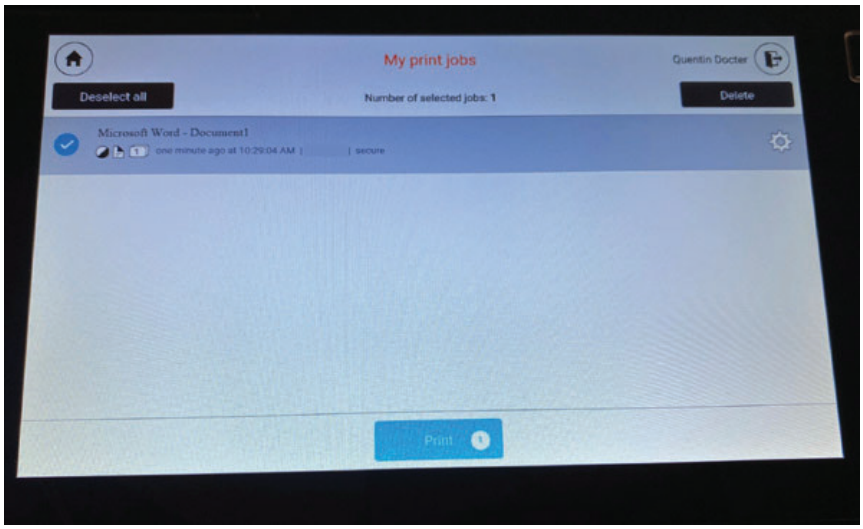

FIGURE 5.39 Xerox print log

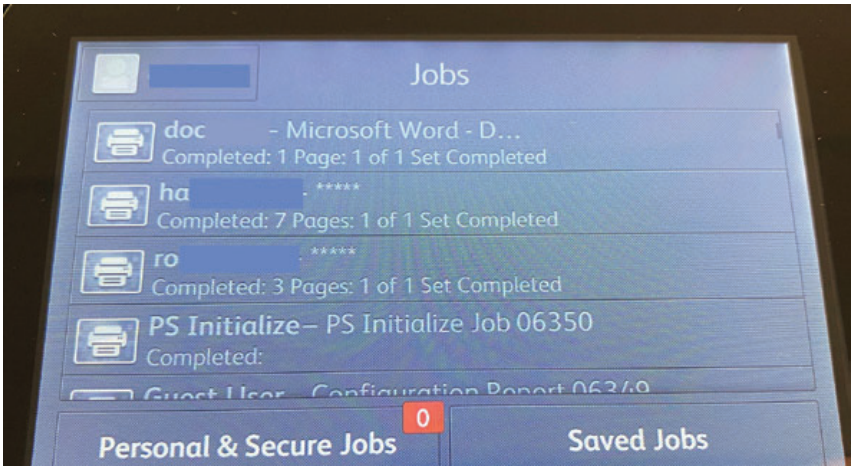
Jobs
doc - Microsoft Word - D...
Completed: 1 Page: 1 of 1 Set Completed
ha *****
Completed: 7 Pages: 1 of 1 Set Completed
ro *****
Completed: 3 Pages: 1 of 1 Set Completed
PS Initialize– PS Initialize Job 06350
Completed:
0
Personal & Secure Jobs
Saved Jobs

Other printers will integrate logging software into the operating system's standard logging utilities. For example, some HP printers will install an audit log into Windows Event Viewer. (We will cover Event Viewer in Volume 2 Chapter 2, "Windows Configuration.") Third-party audit software is also available for use.

Another security issue is storing print jobs on print servers. Print jobs are cached on the print server (it's called *hard drive caching*), and they can be saved indefinitely. We looked at the check box to keep printed documents in Figure 5.31. If you are configuring a printer for a public area, do not select this option. If you're printing to a public printer (such as at

a hotel), and you can see the printer properties, check to make sure that this option is not selected. That will help ensure that your files will not be cached on the hotel computer's hard drive.

Network Scan Services

Scanning is really the opposite of printing. Instead of printing to get electronic information onto paper, scanning takes printed information and stores it electronically. In the office environment, the two most popular types of scanners are flatbed scanners and automatic document feeder (ADF) scanners.

To use a *flatbed scanner*, simply open the lid and lay the document to be scanned on the scanner glass, aligning it to the proper corner as shown in Figure 5.40. There will be an icon on the scanning bed that shows where the document should go. Close the lid and press Scan on the touchscreen, and scanning will begin.

FIGURE 5.40 A flatbed scanner

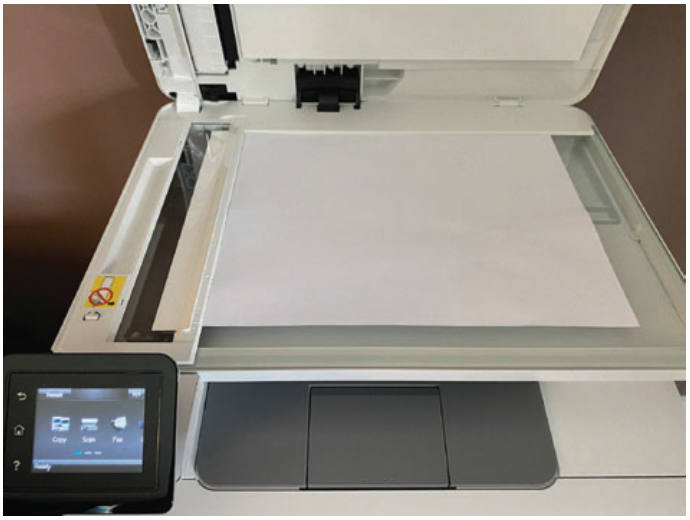

If you have more than one document to scan, using a flatbed scanner can be a real pain. Fortunately, some scanners have an attachment called an *automatic document feeder (ADF)* that lets you scan multiple pieces of paper in one job. Figure 5.41 shows an example with paper loaded in it. It's common to have ADFs that allow for up to 50 pages to be scanned at once.

Automatic document feeder (ADF), flatbed scanner, and network scan services (email, SMB, cloud services) are part of exam objective 3.7. Understand how they are used.

FIGURE 5.41 Automatic document feeder (ADF) on an MFD

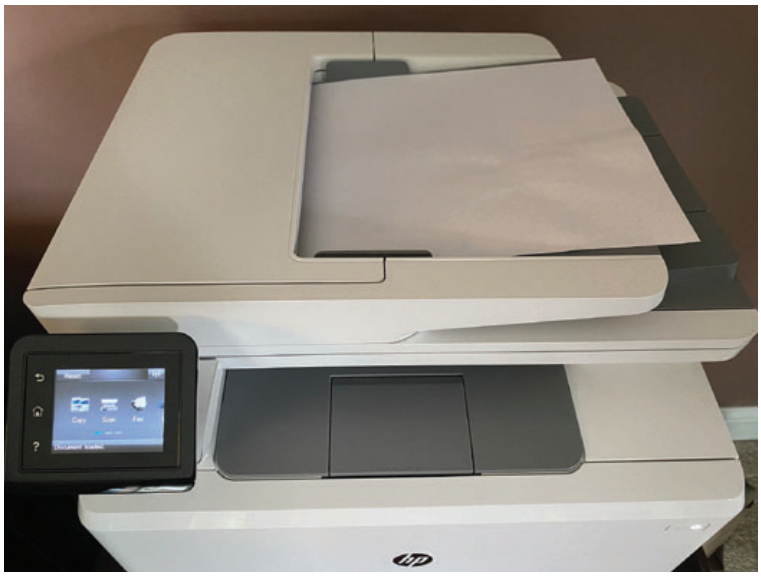

Whenever you scan a document, you need to figure out where to send it. MFDs don't usually have the memory to save images of scans—besides, the point of scanning a document is usually to email it to someone or save it on a hard drive for later retrieval. Let's take a look at three different ways to send or save scanned materials. For all three of these options, it's assumed that the scanner is connected to the network.

Scan to Email Using a *scan to email* option, the scanner will simply email the file, usually in PDF format, to an email recipient. This is probably the most common method that people use when scanning files. Figure 5.42 shows an example of setting up scan to email.

FIGURE 5.42 Scan to email

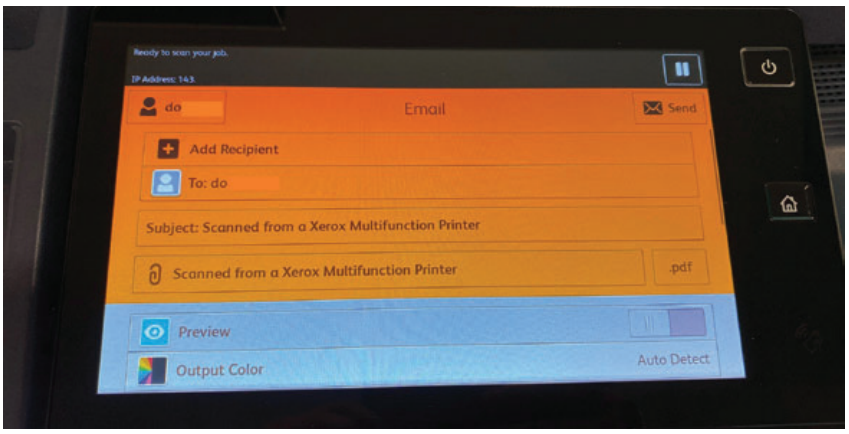

Scan to Folder An alternative option to emailing a scanned file is to save it in a network folder. This is a particularly viable solution if the scanned file is too large to be emailed.

The protocol the printer uses to transport the file from itself to the network folder is called *Server Message Block (SMB)*. In addition, the administrator must set up the MFD to support SMB scanning, and the recipient folder needs to be properly shared and secured too. Performing the scan from the MFD is done via a screen similar to the one shown in Figure 5.42, except instead of selecting an email recipient, you would navigate to the folder where you want to save the file.

Scan to Cloud The final option is to use *cloud services* to save the scan. Various file saving and sharing services support scanning to the cloud, such as Google Drive, OneDrive, Box, and Dropbox, among others. In addition, some printer manufacturers such as HP and Epson offer scan to cloud services. The exact process to follow depends on which cloud service you use to save the file. Exercise 5.5 has you "scan" a document to Google Drive—and it doesn't require you to have an actual scanner! Of course, you do need an Android device to make it work.

EXERCISE 5.5

Scanning a Document to Google Drive

This exercise requires the use of an Android phone or tablet.

1. On the Android device, open the Google Drive app.
2. In the bottom-right corner, tap New (like in Figure 5.43). It should have a plus sign next to it.

FIGURE 5.43 New button in Google Drive

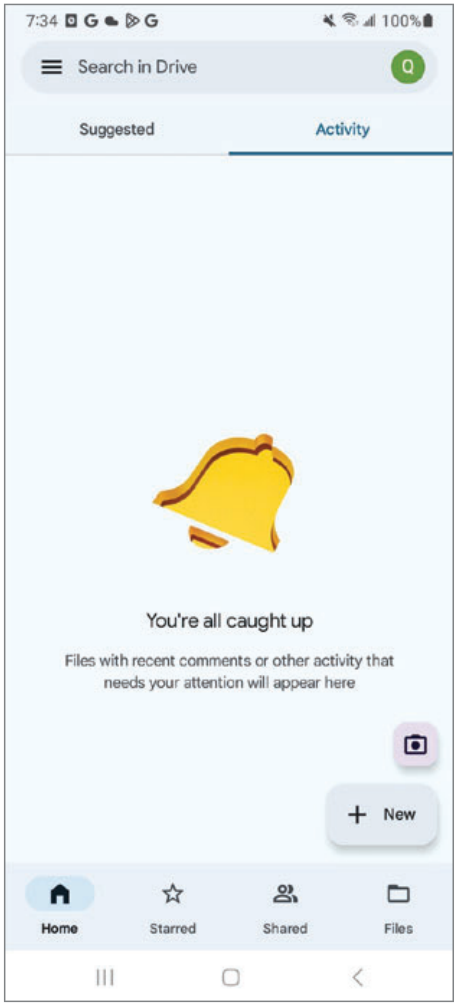

3. Tap Scan. It should have a camera icon next to it.

You might need to give Drive permission to use the camera.

4. Take a photo of the object you want to scan.

5. Tap Done to save the scanned document to Google Drive.

To make the process even shorter, you can install the Drive scan widget on your Home screen.

Performing Appropriate Printer Maintenance

Considering the amount of work they do, printers last a pretty long time. Some printers handle over 100,000 pages per month, yet they're usually pretty reliable devices. You can help your printers live long and fulfilling lives by performing the right maintenance, and smoothly running printers always make your officemates happy. After all, going to get your print job from the printer and discovering that the printer is in the shop is a very frustrating experience!

Regardless of the type of printer you use, giving it a regular checkup is a good idea. You're probably familiar with some of the activities that fall under maintenance, such as replacing paper, ink, or toner cartridges. We'll look at those as well as some additional, more involved maintenance procedures.

Performing appropriate printer maintenance is a major part of Exam Objective 3.8. Be sure to understand how to do the following:

- **Laser**: Replace toner, apply maintenance kit, calibrate, and clean.
- **Inkjet**: Clean printheads, replace cartridges, calibrate, and clear jams.
- **Thermal**: Replace paper, clean heating element, and remove debris.
- **Impact**: Replace ribbon, printhead, and paper.

Replacing Consumables

To maintain a printer properly, you need to replace consumables such as toner or ink cartridges, assemblies, filters, and rollers on occasion. Trying to cut costs by buying cheaper supplies rarely pays off.

Whenever purchasing supplies for your printer, always get supplies from the manufacturer or from an authorized reseller. This way, you'll be sure that the parts are of high quality. Using unauthorized parts can damage your printer and possibly void your warranty.

Printer Paper

Most people don't give much thought to the kind of paper they use in their printers. It's a factor that can have a tremendous effect on the quality of the hard-copy printout, however, and the topic is more complex than people think. For example, the wrong paper can cause frequent paper jams and possibly even damage components.

Several aspects of paper can be measured; each gives an indication as to the paper's quality. The first factor is *composition*. Paper is made from a variety of substances. Paper used to be made from cotton and was called rag stock. It can also be made from wood pulp, which is cheaper. Most paper today is made from the latter or a combination of the two.

Another aspect of paper is the property known as *basis weight* (or simply *weight*). The weight of a particular type of paper is the actual weight, in pounds, of a ream (500 sheets) of the standard size of that paper made of that material. For regular bond paper, that size is 17 × 22.

The final paper property we'll discuss is the *caliper* (or thickness) of an individual sheet of paper. If the paper is too thick, it may jam in feed mechanisms that have several curves in the paper path. On the other hand, a paper that's too thin may not feed at all.

These are just three of the categories we use to judge the quality of paper. Because there are so many different types and brands of printers as well as paper, it would be impossible to give the specifications for the "perfect" paper. However, the documentation for any printer will give specifications for the paper that should be used in that printer.

Many impact printers need to use special paper that has tractor feed perforations on the side, or they will not work properly. When replacing tractor feed paper, it's very easy to get it misaligned, and it will feed crookedly and ultimately jam the printer. Similarly, thermal printers also require special paper that needs to be loaded properly. In many cases, if you load it upside down, the unit will not produce images. By comparison, adding paper to a laser or inkjet printer is usually very easy.

For best results, store paper in an area where it will not get wet or be exposed to excessive humidity.

Ink and Toner

The area in which using recommended supplies is the biggest concern is ink and toner cartridges. Using the wrong ink or toner supplies is the easiest way to ruin a perfectly good printer.

Dot-matrix printers use a cloth or polyester ribbon soaked in ink and coiled up inside a plastic case. This assembly is called a *printer ribbon* (or *ribbon cartridge*). Once the ribbon has run out of ink, it must be discarded and replaced. Ribbon cartridges are developed closely with their respective printers. For this reason, ribbons should be purchased from the same manufacturer as the printer. The wrong ribbon could jam in the printer as well as cause quality problems.

It is possible to re-ink a ribbon. Some vendors sell a bottle of ink solution that can be poured into the plastic casing, where the cloth ribbon will soak up the solution. This can be a messy process, and you should do this only if the manufacturer recommends it.

Inkjet cartridges have a liquid ink reservoir. The ink in these cartridges is sealed inside. Once the ink runs out, the cartridge must be removed and discarded. A new, full cartridge is installed in its place. Because the ink cartridge contains the printing mechanism as well as ink, it's like getting a new printer every time you replace the ink cartridge.

In some inkjet printers, the ink cartridge and the printhead are in separate assemblies. This way, the ink can be replaced when it runs out, and the printhead can be used several times. This approach works fine if the printer is designed to work this way. However, some people think that they can do this on their integrated cartridge/printhead system, using special ink cartridge refill kits. These kits consist of a syringe filled with ink and a long needle. The needle is used to puncture the top of an empty ink cartridge, and the syringe is then used to refill the reservoir.

Do not use ink cartridge refill kits! There are several problems with these kits (the ones you see advertised with a syringe and a needle). First, the kits don't use the same kind of ink that was originally in the ink cartridges. The new ink may be thinner, causing it to run out or not print properly. Also, the printhead is often supposed to be replaced around this same time. Refilling the cartridge doesn't replace the printhead, so you'll have print-quality problems. Finally, the hole the syringe leaves cannot be plugged and may allow ink to leak out. These problems can happen with do-it-yourself kits as well as with cartridges refilled by office supply stores or private printer supply sellers. Here's the bottom line: *Buy new ink cartridges from the printer manufacturer*. Yes, they are a bit more expensive, but in the long run you will save money because you won't have any of the problems described here.

The final type of consumable is toner. Each model of laser printer uses a specific toner cartridge. You should check the printer's manual to see which toner cartridge your printer needs. Many businesses will recycle your toner or ink cartridges for you, refill them, and sell them back to you at a discount. Don't buy them. While some businesses that perform this "service" are more legitimate than others, using recycled parts is more dangerous to your hardware than using new parts. The reason for this is that refilled cartridges are more likely to break or leak than new parts, and this leakage could cause extensive damage to the inside of your printer. And again, using secondhand parts can void your warranty, so you're left with a broken printer that you have to pay for. Avoid problems like this by buying new parts.

Real World Scenario

Think Before You Refill

Just as with ink cartridges, you should always buy the exact model of toner cartridge recommended by the manufacturer. The toner cartridges have been designed specifically for a particular model. Additionally, *never* refill toner cartridges, for most of the same reasons we don't recommend refilling ink cartridges. The printout quality will be poor, and the fact that you're just refilling the toner means that you might *not* be replacing the photosensitive drum (which is usually inside the cartridge), and the drum might *need* to be replaced.

Simply replacing refilled toner cartridges with proper, name-brand toner cartridges has solved most laser printer quality problems that we have run across. We keep recommending the right ones, but clients keep coming back with the refilled ones. The result is that we take our clients' money to solve their print-quality problems when all it involves is a toner cartridge, our (usually repeat) advice to buy the proper cartridge next time, and the obligatory minimum charge for a half hour of labor (even though the job of replacing the cartridge takes all of 5 minutes).

Always properly recycle your used ink and toner cartridges. Just don't buy recycled cartridges!

Performing Scheduled Maintenance

When shopping for a printer, one of the characteristics you should look for is the printer's capacity, which is often quoted in monthly volume. This is particularly important if the printer will be serving in a high-load capacity. Every printer needs periodic maintenance, but printers that can handle a lot of traffic typically need it less frequently. Check the printer specifications to see how often scheduled maintenance is suggested. Never, ever fall behind on performing scheduled maintenance on a printer.

Many laser printers have LCD displays that provide useful information, such as error messages or notices that you need to replace a toner cartridge. The LCD display will also tell you when the printer needs scheduled maintenance. How does it know? Printers keep track of the number of pages they print, and when the page limit is reached, they display a message, usually something simple like *Perform user maintenance*. The printer will still print, but you should perform the maintenance.

Being the astute technician that you are, you clean the printer with the recommended cleaning kit or install the maintenance kit that you purchased from the manufacturer. Now, how do you get the maintenance message to go away? Reset the page count using a menu option. For example, on many HP laser printers, you press the Menu button until you get to the Configuration menu. Once there, you press the Item key until the display shows *Service Message = ON*. Then press the plus key (+) to change the message to *Service Message = OFF*. Bring the printer back online, and you're ready to go.

When performing maintenance on an impact printer, always carefully inspect the printhead for damage. Replace damaged printheads with authorized parts from the manufacturer.

Performing routine maintenance will keep the printer clean, make it last longer, and help prevent annoying paper jams.

Using Cleaning Solutions

With all of the ink or toner they use, printers get dirty. If printers get too dirty or if the printheads get dirty, you'll notice print problems. No one wants this to happen.

Most printers have a self-cleaning utility that is activated through a menu option or by pressing a combination of buttons on the printer itself. It's recommended that you run the cleaning sequence every time you replace the toner or ink cartridges. If you experience print-quality problems, such as lines in the output, run the cleaning routine.

Sometimes, the self-cleaning routines aren't enough to clear up the problem. If you are having print-quality issues, you might want to consider purchasing a cleaning or maintenance kit, which frequently comes with a cleaning solution.

Cleaning kits are often designed for one specific type of printer and should be used only on that type of printer. For example, don't apply an inkjet cleaning solution to a laser printer.

Each cleaning kit comes with its own instructions for use. Exercise 5.6 walks you through the steps of using an inkjet cleaning solution. Note that the steps for your printer might differ slightly; please consult your manual for specific instructions. After using a cleaning kit on a laser or inkjet printer, it's best to perform a calibration per the printer's instructions.

EXERCISE 5.6

Using an Inkjet Cleaning Solution

1. Power on the printer, and open the top cover to expose the area containing the print cartridges.
2. Initiate a self-cleaning cycle. When the printhead moves from its resting place, pull the AC power plug. This lets you freely move the printheads without damaging them.
3. Locate the sponge pads on which to apply the cleaning solution.

   They'll be in the area where the printheads normally park.
4. Use a cotton swab or paper towel to gently soak up any excess ink in the pads.
5. Using the supplied syringe, apply the cleaning solution to the sponge pads until they are saturated.
6. Plug the printer back into the wall outlet, and turn it on. The printheads will park themselves.
7. Turn the printer back off. Let the solution sit for at least 3 hours.
8. Power the printer back on, and run three printer cleaning cycles.
9. Print a nozzle check pattern (or a test page) after each cleaning cycle to monitor the cleaning progress.

That should take care of it. If not, refer to your printer's manual for more instructions.

Thermal printers require special attention because they contain a heating element. Always unplug the device and ensure that it's cooled off before trying to clean it. Thermal printer cleaning cards, cleaning pens, and kits are widely available in the marketplace. If you need to remove any debris (from any printer), use compressed air or a specialized computer vacuum.

When using compressed air, take care to not just blow the dust or debris all over the place or onto other components. Have a plan in place to contain it all before squeezing the trigger!

Ensuring a Suitable Environment

Printers won't complain if the weather outside is too hot or too cold, but they are susceptible to environmental issues. Here are some things to watch out for in your printer's environment:

Heat Laser printers and thermal printers can generate a lot of heat. Because of this, ensure that your printer is in a well-ventilated area. Resist the temptation to put the printer in the little cubbyhole in your desk. Overheating will reduce the shelf life of your printer.

Humidity High humidity can cause printer paper to stick together. Sticky paper leads to paper jams. Humidity over 80 or 90 percent can cause issues.

Light The laser printer's toner cartridge contains a photosensitive drum. Exposing that drum to light could ruin the drum. While the drum is encased in plastic, it's best to avoid exposing the printer or toner cartridges to extreme light sources. Under no circumstance should you open the toner cartridge, unless you're ready to get rid of it as well as clean up a big mess.

Ozone Laser printers that use corona wires produce ozone as a by-product of the printing process. In offices, ozone can cause respiratory problems in small concentrations, and it can be seriously dangerous to people in large amounts. Ozone is also a very effective oxidizer and can cause damage to printer components.

Fortunately, laser printers don't produce large amounts of ozone, and most laser printers have an ozone filter. Ozone is another reason to ensure that your printer area has good ventilation. Also, replace the ozone filter periodically; check your printer's manual for recommendations on when to do this.

Ammonia A printer doesn't produce ammonia, but it is contained in many cleaning products. Ammonia can greatly reduce the printer's ability to neutralize ozone and can cause permanent damage to toner cartridges. It's best to avoid using ammonia-based cleaners near laser printers.

Summary

In this chapter, we discussed how different types of printers and multifunction devices work as well as the most common methods of connecting them to computers. You learned how computers use page-description languages to format data before they send it to printers and drivers to talk to them. You also learned about the various types of consumable supplies and how they relate to each type of printer.

The most basic category of printer currently in use is the impact printer. Impact printers form images by striking something against a ribbon, which in turn makes a mark on the paper. You learned how these printers work and the service concepts associated with them.

One of the most popular types of printer today is the inkjet printer, so named because of the mechanism used to put ink on the paper.

The most complex type of printer is the laser printer. You learned about the steps in the electrophotographic (EP) imaging process, the process that explains how laser printers print. We also explained the various components that make up this printer and how they work together.

The last printer type we looked at was thermal, which uses special heat sensitive thermal paper that's heat sensitive to create output. Be sure to store the special thermal paper in a cool place or it could get ruined.

You then learned about the interfaces used to connect printers to PCs and how to install and share a printer. Common interfaces include USB, Ethernet, and wireless. Proper steps include connecting the device, installing the driver, configuring options, validating application and operating system compatibility, and educating users on how to use the device. Installing the device is the first step, but you're not done until you ensure that it works properly and that users know how to access it.

Installing network printers usually involves a few more steps than are needed to install local printers, and the device is connected to the network instead of to a host. Networked MFDs are often used for scan services, such as scanning to email, SMB, and the cloud. Security becomes critical here as well, so you should be familiar with user authentication, badging, secured prints, and audit logs.

Finally, we looked at how to perform printer maintenance, including the importance of using recommended supplies.

Exam Essentials

Know the differences between types of printer technologies (for example, laser, inkjet, thermal, impact). Laser printers use a laser and toner to create the page. Inkjet printers spray ink onto the page. Thermal printers use heat to form the characters on the page. Impact printers use a mechanical device to strike a ribbon, thus forming an image on the page.

Know the three most common ways to connect a printer. The methods are USB, Ethernet, and wireless.

Be familiar with printer configuration settings. Know how duplex, orientation, tray settings, and print quality are configured.

For networked printers, understand security and scan services options. Security can include user authentication and badging, audit logs, and secured prints. Network scan services include scan to email, scan to a folder (using the SMB protocol), and scan to cloud.

Understand what print servers and printer shares are. A shared printer is one that's made available to users across the network. The print server is the computer or system that hosts the shared printer.

Know how to install and configure printers. The basic procedure is as follows:

1. Find a suitable location for the printer and properly unbox the device.
2. Attach the device using a local or network port and connect the power.
3. Install and update the device driver and calibrate the device.
4. Configure options and default settings.
5. Print a test page.
6. Verify compatibility with the operating system and applications.
7. Educate users about basic functionality.

Know the seven steps in the laser imaging process. The seven steps are processing, charging, exposing, developing, transferring, fusing, and cleaning.

Know the key parts in a laser printer and appropriate maintenance procedures. Key parts are the imaging drum, fuser assembly, transfer belt, transfer roller, pickup rollers, separation pads, and duplexing assembly. Maintenance includes replacing toner, applying a maintenance kit, calibrating, and cleaning.

Know the key parts in an inkjet printer and appropriate maintenance procedures. Inkjet parts include the ink cartridge, printhead, roller, feeder, duplexing assembly, and carriage belt. Maintenance items include cleaning printheads, replacing cartridges, calibrating, and clearing paper jams.

Know the key components in a thermal printer and appropriate maintenance procedures. The feed assembly and heating element are important thermal printer parts. The paper is also important here because it's special heat-sensitive thermal paper. Maintenance includes replacing paper, cleaning the heating element, and removing debris.

Know the key parts in an impact printer and appropriate maintenance procedures. Impact printer parts to know include the printhead, ribbon, tractor feed (and special tractor feed paper), and impact or multipart paper. Maintenance includes replacing the ribbon, printhead, and paper.

Understand the importance of using recommended supplies. Using consumables (paper, ink, toner) that are recommended for your printer is important. Using bad supplies could ruin your printer and void your warranty.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. A user has a printer attached to their computer. They configured the printer so other users on the network can print to it. What does that make their computer?
   - **A.** Print server
   - **B.** Printer share
   - **C.** Network scan service
   - **D.** ADF

2. Which types of printers are referred to as page printers because they receive their print job instructions one page at a time? (Choose two.)
   - **A.** Impact
   - **B.** Inkjet
   - **C.** Laser
   - **D.** Thermal

3. Which of the following is not an advantage of a Universal Serial Bus (USB) printer interface?
   - **A.** It has a higher transfer rate than a serial connection.
   - **B.** It has a higher transfer rate than a parallel connection.
   - **C.** It automatically recognizes new devices.
   - **D.** It allows the printer to communicate with networks, servers, and workstations.

4. You have a finance manager who needs to print contracts in triplicate. Which type of printers are best for printing on multipart forms?
   - **A.** Inkjet printers
   - **B.** Laser printers
   - **C.** Thermal printers
   - **D.** Dot-matrix printers

5. Which step in the EP imaging process uses a laser to discharge selected areas of the photosensitive drum, thus forming an image on the drum?
   - **A.** Exposing
   - **B.** Transferring
   - **C.** Developing
   - **D.** Cleaning

6. Which of the following is a type of device that lets users scan in 10-page double-sided documents all at once?
   - **A.** SMB
   - **B.** MFD
   - **C.** ADF
   - **D.** Flatbed

7. Recently, employees have printed off unauthorized materials which may contain prohibited hate speech. Which of the following can be used to see which users printed specific jobs?
   - **A.** Audit logs
   - **B.** Badging
   - **C.** User authentication
   - **D.** Secured prints

8. Which device in an inkjet printer contains the printhead?
   - **A.** Toner cartridge
   - **B.** Ink cartridge
   - **C.** Daisy wheel
   - **D.** Paper tray

9. What is the correct order of the steps in the EP imaging process?
   - **A.** Developing, exposing, transferring, fusing, charging, cleaning, processing
   - **B.** Charging, processing, exposing, developing, transferring, fusing, cleaning
   - **C.** Processing, transferring, exposing, developing, charging, cleaning, fusing
   - **D.** Processing, charging, exposing, developing, transferring, fusing, cleaning

10. The network administrator has set up security on office MFDs that requires users to authenticate before printing. Which of the following are valid methods for users to authenticate to the MFD? (Choose two.)
    - **A.** Use an authentication log.
    - **B.** Enable secured print.
    - **C.** Type their username and password on the MFD's touchscreen.
    - **D.** Use SMB.
    - **E.** Use their company ID for badging.

11. What is typically included in the EP laser printer toner cartridge? (Choose three.)
    - **A.** Toner
    - **B.** Print drum
    - **C.** Laser
    - **D.** Cleaning blade

12. What happens during the developing stage of laser printing?
    - **A.** An electrostatic charge is applied to the drum to attract toner particles.
    - **B.** Heat is applied to the paper to melt the toner.
    - **C.** The laser creates an image of the page on the drum.
    - **D.** An electrostatic charge is applied to the paper to attract toner particles.

13. Which of the following are possible interfaces for printers? (Choose three.)

A. Parallel

B. SATA

C. USB

D. Network

14. You have just installed a new printer, but it prints only garbled text. Which of the following is most likely the problem?

A. Wrong IP address

B. Worn printhead

C. Incorrect printer drivers

D. Unsupported printer

15. Which printer contains a wheel that looks like a flower with raised letters and symbols on each "petal"?

A. Inkjet printer

B. Daisy-wheel printer

C. Dot-matrix printer

D. Laser printer

16. What part of a laser printer supplies the voltages for charging and transferring corona assemblies?

A. High-voltage power supply (HVPS)

B. DC power supply (DCPS)

C. Controller circuitry

D. Transfer corona

17. Which printer part gets the toner from the photosensitive drum onto the paper?

A. Laser-scanning assembly

B. Fusing assembly

C. Corona assembly

D. Drum

18. Which step in the laser printer printing process occurs immediately after the exposing phase?

A. Charging

B. Fusing

C. Transferring

D. Developing

19. Which laser printer component permanently presses the toner into the paper?

A. Transfer corona

B. Fuser assembly

C. Printer controller circuitry

D. Paper transport assembly

20. Which of the following most accurately describes how to obtain a firmware upgrade for your laser printer?

A. Download the firmware upgrade for free from the manufacturer's website.

B. Pay to download the firmware upgrade from the manufacturer's website.

C. Have a certified laser printer technician come to your site and install a new firmware chip.

D. Contact the manufacturer of the printer, and they will send you the firmware upgrade on an SSD.

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answer compares to the authors', refer to Appendix B.

Your network has several inkjet printers in use. A user is complaining that their documents are consistently printing with extra smudges along the lines of print on one of them. What steps would you take to clean the printer?

Chapter 6

Networking Fundamentals

THE FOLLOWING COMPTIA A+ 220-1201 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

✓ **2.0 Networking**

✓ **2.5 Compare and contrast common networking hardware devices.**

- Routers
- Switches
  - Managed
  - Unmanaged
- Access points
- Patch panel
- Firewall
- Power over Ethernet (PoE)
  - Injectors
  - Switch
  - PoE standards
- Cable modem
- Digital subscriber line (DSL)
- Optical network terminal (ONT)
- Network interface card (NIC)
  - Physical media access control (MAC) address

✓ **2.7 Compare and contrast internet connection types, network types, and their characteristics.**

- Network types
- LAN

- Wide area network (WAN)
- Personal area network (PAN)
- Metropolitan area network (MAN)
- Storage area network (SAN)
- Wireless local area network (WLAN)

✓ **3.0 Hardware**

✓ **3.2 Summarize basic cable types and their connectors, features, and purposes.**

- Network cables
  - Copper
    - Categories
      - T568A/T568B standards
    - Coaxial
    - Shielded twisted pair
      - Direct burial
    - Unshielded twisted pair
  - Plenum-rated
  - Optical
- Single-mode
- Multimode
- Connector types
  - RJ11
  - RJ45
  - F-type
  - Straight tip (ST)
  - Subscriber connector (SC)
  - Lucent connector (LC)
  - Punchdown block

Looking around most homes or offices today, it's hard to imagine a world without networks. Nearly every place of business has some sort of network. Wireless home networks have exploded in popularity in the last decade, and it seems that everywhere you go, you can see a dozen wireless networks from your smartphone, smartwatch, tablet, or laptop.

It didn't used to be that way. Even when not thinking about networks, we're still likely connected to one via the ubiquitous Internet-enabled smartphones in our pockets. We take for granted a lot of what we have gained in technology over the past few years, much less the past several decades.

Thirty years ago, if you wanted to send a memo to everyone in your company, you had to use a photocopier and interoffice mail. Delivery to a remote office could take days. Today, one mistaken click of the Reply All button can result in instantaneous embarrassment. Email is an example of one form of communication that became available with the introduction and growth of networks.

This chapter focuses on the basic concepts of how a network works, including the way it sends information, the hardware used, and the common types of networks you might encounter. It used to be that in order to be a PC technician, you needed to focus on only one individual (but large) computer at a time. In today's environment, though, you will in all likelihood need to understand combinations of hardware, software, and network infrastructure in order to be successful.

If the material in this chapter interests you, you might consider studying for, and eventually taking, CompTIA's Network+ exam. It is a non–company-specific networking certification similar to A+ but for network-related topics. You can study for it using Sybex's *CompTIA Network+ Study Guide*, by Todd Lammle and Jon Buhagiar (May 7, 2024, Sybex), available at your favorite online bookseller.

Understanding Networking Principles

Stand-alone personal computers, first introduced in the late 1970s, gave users the ability to create documents, spreadsheets, and other types of data and save them for future use. For the small-business user or home-computer enthusiast, this was great. For larger companies, however, it was not enough. Larger companies had greater needs to share information between

offices and sometimes over great distances. Stand-alone computers were insufficient for the following reasons:

- Their small hard drive capacities were insufficient.
- To print, each computer required a printer attached locally.
- Sharing documents was cumbersome. People grew tired of having to save to a floppy and then take that disk to the recipient. (This procedure was called *sneakernet*.)
- There was no email. Instead, there was interoffice mail, which was slow and sometimes unreliable. Imagine snail mail without post office efficiency.

To address these problems, networks were born. A *network* links two or more computers together to communicate and share resources. Their success was a revelation to the computer industry as well as to businesses. Now departments could be linked internally to offer better performance and increase efficiency.

The term *networking* can be used in a social or business context, where people come together and exchange names for future contact and access to more resources. The same is true with a computer network. A computer network enables computers to link to each other's resources. For example, in a network, every computer does not need a printer connected locally in order to print. Instead, you can connect a printer to one computer, or you can connect it directly to the network and allow all the other computers to access it. Because they allow users to share resources, networks can increase productivity as well as decrease cash outlay for new hardware and software.

In many cases, networking today has become a relatively simple plug-and-play process. Wireless network cards can automatically detect and join networks, and then you're seconds away from surfing the web or sending email. Of course, not all networks are that simple. Getting your network running may require a lot of configuration, and one messed-up setting can cause the whole thing to fail.

To best configure your network, there is a lot of information you should understand about *how* networks work. This section covers the fundamentals, and armed with this information, you can then move on to how to make it work *right*.

Network Types

One of the ways to think about how networks are structured is to categorize them by network type. Some networks are small in scale whereas others span the globe. Some are designed to be wireless only, whereas others are designed specifically for storage. Understanding the basic structure of the network can often help you solve a problem. There are six different types of networks you need to be familiar with, and we'll cover them here:

- Local area network (LAN)
- Wide area network (WAN)
- Personal area network (PAN)
- Metropolitan area network (MAN)

- Storage area network (SAN)
- Wireless local area network (WLAN)

Exam Objective 2.7, "Compare and contrast internet connection types, network types, and their characteristics," specifies that you should know all six network types in the bulleted list. It also calls out Internet connection types. We will cover those in Volume 1 Chapter 9, "Wireless and SOHO Networking," where we show you how to set up a network from start to finish.

Local Area Networks (LANs)

The 1970s brought us the minicomputer, which was a smaller version of large mainframe computers. Whereas the mainframe used *centralized processing* (all programs ran on the same computer), the minicomputer used *distributed processing* to access programs across other computers. As depicted in Figure 6.1, distributed processing allows a user at one computer to use a program on another computer as a *backend* to process and store information. The user's computer is the *frontend*, where data entry and minor processing functions are performed. This arrangement allowed programs to be distributed across computers rather than be centralized. This was also the first time network cables rather than phone lines were used to connect computers.

FIGURE 6.1 Distributed processing

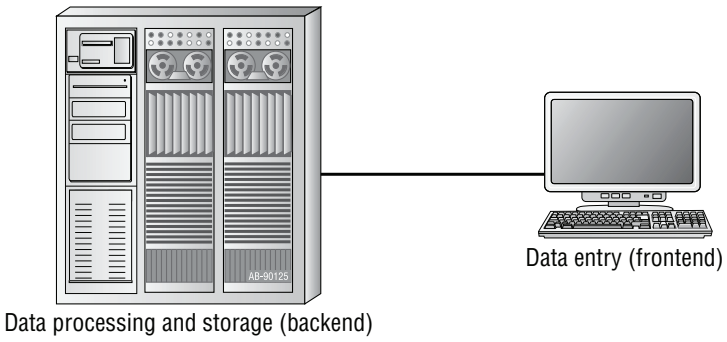

By the 1980s, offices were beginning to buy PCs in large numbers. Portables (one step toward laptops, but they were way too big to sit on anyone's lap) were also introduced, allowing computing to become mobile. Neither PCs nor portables, however, were efficient in sharing information. As timeliness and security became more important, floppy disks were just not cutting it. Offices needed to find a way to implement a better means to share and access resources. This led to the introduction of the first type of PC *local area network (LAN)*: ShareNet by Novell, which had both hardware and software components. LANs simply link computers in order to share resources within a closed environment, such as a single office or building. The first simple LANs were constructed a lot like the LAN shown in Figure 6.2.

FIGURE 6.2 A simple LAN

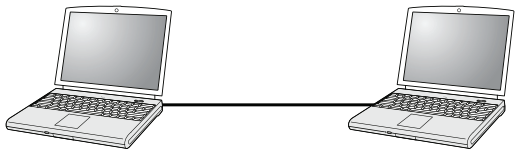

After the introduction of ShareNet, more LANs sprouted. The earliest LANs could not cover large distances. Most of them could only stretch across a single floor of the office and could support no more than 30 computers. Furthermore, they were still very rudimentary and only a few software programs supported them. The first software programs that ran on a LAN were not capable of being used by more than one user at a time. (This constraint was known as *file locking*.) Nowadays, multiple users often concurrently access a program or file. Most of the time, the only limitations will be restrictions at the record level if two users are trying to modify a database record at the same time.

Wide Area Networks (WANs)

By the late 1980s, networks were expanding to cover large geographical areas and were supporting thousands of users. The concept of a *wide area network (WAN)* was born. WANs were first implemented with mainframes at massive government expense, but started attracting PC users as networks went to this new level. Employees of businesses with offices across the country communicated as though they were only desks apart. Soon the whole world saw a change in the way of doing business, across not only a few miles but across countries. Whereas LANs are limited to single buildings, WANs can span buildings, states, countries, and even continental boundaries. Figure 6.3 shows an example of a simple WAN.

FIGURE 6.3 A simple WAN

Generally speaking, it's safe to think of a WAN as multiple dispersed LANs connected together. Historically, only larger corporations used WANs, but many smaller companies with remote locations now use them as well. The networks of today and tomorrow are no longer limited by the inability of LANs to cover distance and handle mobility. WANs play an important role in the future development of corporate networks worldwide.

Personal Area Networks (PANs)

In moving from LANs to WANs, we increased the scope. Going the other way, a *personal area network (PAN)* is going to be much smaller in scale. The term PAN is most commonly used in reference to Bluetooth networks. In 1998, a consortium of companies formed the Bluetooth Special Interest Group (SIG) and formally adopted the name *Bluetooth* for its technology. The name comes from a tenth-century Danish king named Harald Blåtand, known as Harold Bluetooth in English. (One can only imagine how he got that name.) King Blåtand had successfully unified warring factions in the areas of Norway, Sweden, and Denmark. The makers of Bluetooth were trying to unite disparate technology industries, namely computing, mobile communications, and the auto industry.

Although the most common use of a PAN is in association with Bluetooth, a PAN can also be created with other technologies, such as infrared.

Current membership in the Bluetooth SIG includes Microsoft, Intel, Apple, IBM, Toshiba, and several cell phone manufacturers. The technical specification IEEE 802.15.1 describes a *wireless personal area network (WPAN)* based on Bluetooth version 1.1.

The Institute of Electrical and Electronics Engineers (IEEE, pronounced "eye triple E") is the world's largest technical professional organization. It's a company-agnostic group that is widely recognized as the leader in the development of industry standards for electrical and electronic technologies. IEEE is referenced everywhere in computer networking documentation—you will see its name pop up a lot.

The first Bluetooth device on the market was an Ericsson headset and cell phone adapter, which arrived on the scene in 2000. While mobile phones and accessories are still the most common type of Bluetooth device, you will find many more, including wireless keyboards, mice, and printers. Figure 6.4 shows a Bluetooth USB adapter.

We cover Bluetooth in more detail in Volume 1 Chapter 9 when we discuss wireless technologies. Also, if you want to learn more about Bluetooth, you can visit `www.bluetooth.com`.

FIGURE 6.4 Bluetooth USB adapter

One of the defining features of a Bluetooth WPAN is its temporary nature. With traditional Wi-Fi, you need a central communication point, such as a wireless router or access point, to connect more than two devices together. (This is referred to as *infrastructure*.) Bluetooth networks are formed on an ad hoc basis, meaning that whenever two Bluetooth devices get close enough to each other, they can communicate directly with each other—no central communication point is required. This dynamically created network is called a *piconet*. A Bluetooth-enabled device can communicate with up to seven other devices in one piconet. Two or more piconets can be linked together in a *scatternet*. In a scatternet, one or more devices would serve as a bridge between the piconets.

Metropolitan Area Networks (MANs)

For those networks that are larger than a LAN but confined to a relatively small geographical area, there is the term *metropolitan area network (MAN)*. A MAN is generally defined as a network that spans a city or a large campus. For example, if a city decides to install wireless hotspots in various places, that network could be considered a MAN.

One of the questions a lot of people ask is, "Is there really a difference between a MAN and a WAN?" There is definitely some gray area here; in many cases they are virtually identical. Perhaps the biggest difference is who has responsibility for managing the connectivity. In a MAN, a central IT organization, such as the campus or city IT staff, is responsible. In a WAN, it's implied that you will be using publicly available communication lines, and there will be a phone company or other service provider involved.

Storage Area Networks (SANs)

A *storage area network (SAN)* is designed to do exactly what it says, which is to store information. Although a SAN can be implemented a few different ways, imagine a network (or network segment) that holds nothing but networked storage devices, whether they be network-attached storage (NAS) hard drives or servers with lots of disk space dedicated solely to storage. This network won't have client computers or other types of servers on it. It's for storage only. Figure 6.5 shows what a SAN could look like.

Perhaps you're thinking, why would someone create a network solely for storage? It's a great question, and there are several benefits to having a SAN.

FIGURE 6.5 Storage area network (SAN)

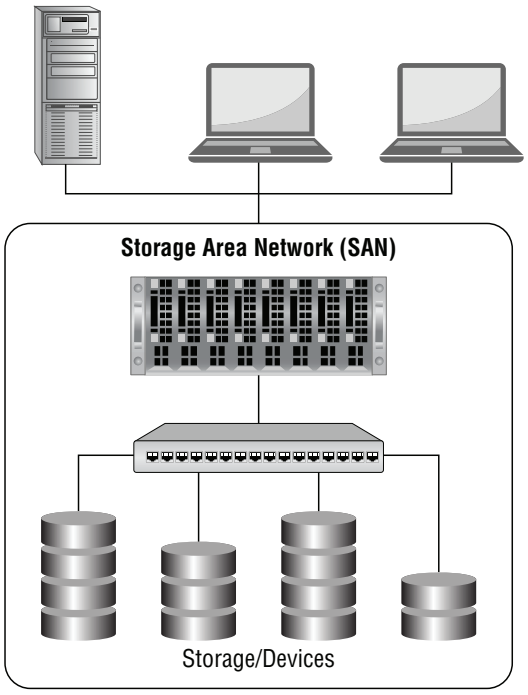

Dedicated SANs relieve network loads. With a SAN, all storage traffic, which may include huge data files or videos, is sent to a specific network or network segment, relieving traffic on other network segments.

SANs offer fast data access. Most SANs use high-speed Fibre Channel connections, which allow for very fast access even for huge files.

SANs are easily expandable. In most cases, it's a matter of connecting a new storage unit, which might even be hot-swappable (removed and replaced without powering down the system), and a few clicks to configure it. Then it's ready to go.

Block-level storage is more efficient. This is getting into the weeds a bit, but most SANs are configured to store and retrieve data in a system called *block storage*. This contrasts with the file-based access systems you're probably used to, such as the ones in Windows and macOS. For anyone who has used a Windows-based or Mac computer, file storage is instantly recognizable. It's based on the concept of a filing cabinet. Inside the filing cabinet are folders, and files are stored within the folders. Each file has a unique name when you include the folders and subfolders it's stored in. For example, `c:\files\doc1.txt` is different from `c:\papers\doc1.txt`. The hierarchical folder structure and the naming scheme of file storage make it relatively easy for humans to navigate. Larger datasets and multiple embedded folders can make it trickier—who here hasn't spent 10 minutes trying to figure out which folder they put that file in?—but it's still pretty straightforward.

With file storage, each file is treated as its own singular entity, regardless of how small or large it is. With block storage, files are split into chunks of data of equal size, assigned a unique identifier, and then stored on the hard drive. Because each piece of data has a unique address, a file structure is not needed. Figure 6.6 illustrates what this looks like.

FIGURE 6.6 Block storage

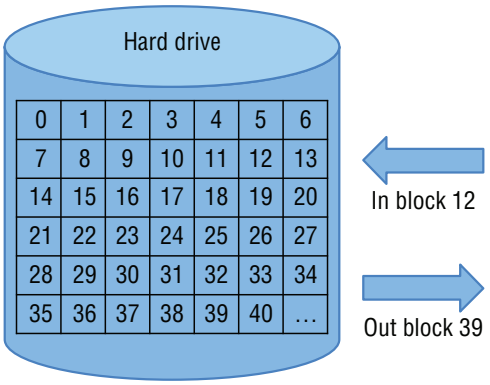

Block storage allows a file to be broken into more manageable chunks rather than being stored as one entity. This allows the operating system to modify one portion of a file without needing to open the entire file. In addition, since data reads and writes are always of the same block size, data transfers are more efficient and therefore faster. Latency (delay) with block storage is lower than with other types of storage.

One of the first common use cases for block storage was for databases, and it remains the best choice for large, structured databases today. Block storage is also used for storage area networks (SANs).

The downsides to SANs are that they are a bit complicated to set up and can be more expensive to run than non-SAN storage solutions. For huge networks that need to get data to large numbers of users, though, they're a good choice.

Wireless Local Area Networks (WLANs)

Wireless networks are everywhere today. If you use your smartphone, tablet, or laptop to look for wireless networks, chances are you will find several. A *wireless local area network (WLAN)* is simply a LAN, but one in which clients connect wirelessly rather than through network cables.

Wireless clients on a network typically access the network through a *wireless access point (WAP)*. The WAP may connect wirelessly to another connectivity device, such as a wireless router, but more likely uses a wired connection to a router or switch. (We'll talk about all of these devices later in the chapter.)

Primary Network Components

Technically speaking, two or more computers connected together constitute a network. But networks are rarely that simple. When you're looking at the devices or resources available on a network, there are three types of components of which you should be aware:

- Servers
- Clients or workstations
- Resources

Servers

Servers come in many shapes and sizes. They are a core component of the network, providing a link to the resources necessary to perform any task. The link that the *server* provides could be to a resource existing on the server itself or to a resource on a client computer. The server is the critical enabler, offering directions to the client computers regarding where to go to get what they need.

Servers offer networks the capability of centralizing the control of resources and security, thereby reducing administrative difficulties. They can be used to distribute processes for balancing the load on computers and can thus increase speed and performance. They can also compartmentalize files for improved reliability. That way, if one server goes down, not all of the files are lost.

Servers can perform several different critical roles on a network. For example, a server that provides files to the users on the network is called a *file server*. Likewise, one that hosts printing services for users is called a *print server*, as you learned in Volume 1 Chapter 5, "Printers and Multifunction Devices." Servers can be used for other tasks as well, such as authentication, remote access services, administration, web hosting, email, and so on. Networks can include multipurpose and single-purpose servers. A multipurpose server can be, for example, both a file server and a print server at the same time. If the server is a single-purpose server, it is a file server only or a print server only. Another distinction we use in categorizing servers is whether they are dedicated or nondedicated:

Dedicated Servers A *dedicated server* is assigned to provide specific applications or services for the network and nothing else. Because a dedicated server specializes in only a few tasks, it requires fewer resources than a nondedicated server might require from the computer that is hosting it. This savings may translate to efficiency and can thus be considered as having a beneficial impact on network performance. A web server is an example of a dedicated server; it is dedicated to the task of serving up web pages and nothing else.

Nondedicated Servers Nondedicated servers are assigned to provide one or more network services *and* local access. A *nondedicated server* is expected to be slightly more flexible in its day-to-day use than a dedicated server. Nondedicated servers can be used

to direct network traffic and perform administrative actions, but they also are often used to serve as a frontend for the administrator to work with other applications or services or to perform services for more than one network. For example, a dedicated web server might serve out one or more websites, whereas a nondedicated web server serves out websites but might also function as a print server on the local network or as the administrator's workstation.

The nondedicated server is not what some would consider a true server, because it can act as a workstation as well as a server. The workgroup server at your office is an example of a nondedicated server. It might be a combination file, print, and email server. Plus, because of its nature, a nondedicated server could also function well in a peer-to-peer environment. It could be used as a workstation in addition to being a file, print, and email server.

We will talk in more depth about server roles in Volume 1 Chapter 8, "Network Services."

Cloud services, which we will talk about in Volume 1 Chapter 11, "Virtualization and Cloud Computing," have changed how servers are perceived, used, and managed. For example, cloud services can include servers, storage, networking hardware, and applications, and where they reside is opaque to the user. Many networks use both dedicated and nondedicated servers as well as cloud services to incorporate the best of all worlds, offering improved network performance with dedicated servers, flexibility with the nondedicated servers, and even more flexibility with the cloud.

Workstations

Workstations are the computers on which the network users do their work, performing activities such as word processing, database design, graphic design, email, and other office or personal tasks. A *workstation* is basically an everyday computer, except for the fact that it is connected to a network that offers additional resources. Workstations can range from diskless computer systems to desktops or laptops. In network terms, workstations are also known as *client computers*. As clients, they are allowed to communicate with the servers in the network to use the network's resources.

It takes several items to make a workstation into a network client. You must install a *network interface card (NIC)*, a special expansion card that allows the PC to talk on a network. You must connect it to a cabling system that connects to other computers (unless your NIC supports wireless networking). And you must install special software, called *client software*, which allows the computer to talk to the servers and request resources from them. Once all this has been accomplished, the computer is "on the network." We'll cover more details on how NICs work and how to configure them in the "Network Interface Cards" section later in this chapter.

Network client software comes with all operating systems today. When you configure your computer to participate in the network, the operating system utilizes this software.

To the client, the server may be nothing more than just another drive letter. However, because it is in a network environment, the client can use the server as a doorway to more storage or more applications or to communicate with other computers or other networks. To users, being on a network changes a few things:

- They can store more information because they can store data on other computers on the network.
- They can share and receive information from other users, perhaps even collaborating on the same document.
- They can use programs that would be too large or complex for their computer to use by itself.
- They can use hardware not attached directly to their computer, such as a printer.

Real World Scenario

Is That a Server or a Workstation?

This is one of the things that author Quentin Docter does when teaching novice technicians. In the room, there will be a standard-looking mini-tower desktop computer. He points to it and asks, "Is that a server or a workstation?" A lot of techs will look at it and say it's a workstation because it is a desktop computer. The real answer is, "It depends."

Although many people have a perception that servers are ultra-fancy, rack-mounted devices, that isn't necessarily true. It's true that servers typically need more powerful hardware than do workstations because of their role on the network, but that doesn't have to be the case. (Granted, having servers that are less powerful than your workstations doesn't make logical sense.) What really differentiates a workstation from a server is what operating system it has installed and what role it plays on the network.

For example, if that system has Windows Server 2025 installed on it, you can be pretty sure that it's a server. If it has Windows 11, it's more than likely going to be a client, but not always. Computers with operating systems such as Windows 11 can be both clients on the network and nondedicated servers, as would be the case if you share your local printer with others on the network.

The moral of the story? Don't assume a computer's role simply by looking at it. You need to understand what is on it and its role on the network to make that determination.

Network Resources

We now have the server to share the resources and the workstation to use them, but what about the resources themselves? A *resource* (as far as the network is concerned) is any item that can be used on a network. Resources can include a broad range of items, but the following items are among the most important:

- Printers and other peripherals
- Disk storage and file access
- Applications

When only a few printers (and all the associated consumables) have to be purchased for the entire office, the costs are dramatically lower than the costs for supplying printers at every workstation.

Networks also give users more storage space to store their files. Client computers can't always handle the overhead involved in storing large files (e.g., database files) because they are already heavily involved in users' day-to-day work activities. By implementing something like a NAS server, a dedicated device can be allocated to store all of the larger files that are used every day, freeing up disk space on client computers. In addition, if users store their files on a server, the administrator can back up the server periodically to ensure that if something happens to a user's files, those files can be recovered.

Files that all users need to access (such as emergency contact lists and company policies) can also be stored on a server. Having one copy of these files in a central location saves disk space, as opposed to storing the files locally on everyone's system.

Applications (programs) no longer need to be on every computer in the office. If the server is capable of handling the overhead that an application requires, the application can reside on the server and be used by workstations through a network connection. Apps can also be cloud-based, which basically means they will reside on a computer somewhere on the Internet. We cover cloud computing in Volume 1 Chapter 11.

The sharing of applications over a network requires a special arrangement with the application vendor, who may wish to set the price of the application according to the number of users who will be using it. The arrangement allowing multiple users to use a single installation of an application is usually part of a *corporate-use license*. We will cover software licensing in Volume 2 Chapter 9, "Safety and Environmental Concerns."

Being on a Network Brings Responsibilities

You are part of a community when you are on a network, which means that you need to take responsibility for your actions. First, a network is only as secure as the users who use it. You cannot randomly delete files or move documents from server to server. You do not own your email, so anyone in your company's management team can choose to

read it. In addition, sending something to the printer does not necessarily mean that it will print immediately—your document may not be the first in line to be printed at the shared printer. Plus, if your workstation has also been set up as a nondedicated server, you cannot turn it off.

Expectations of user behavior are generally set forth in documents such as an acceptable use policy (AUP), service level agreement (SLA), and nondisclosure agreement (NDA). Every organization should have them. Users should be trained on these policies and certify that they are aware of and bound to them.

Network Resource Access

We have discussed two major components of a typical network—servers and workstations—and we've also talked briefly about network resources. Let's dive a bit deeper into how those resources are accessed on a network.

There are generally two resource access models: peer-to-peer and client-server. It is important to choose the appropriate model. How do you decide which type of resource model is needed? You must first think about the following questions:

- What is the size of the organization?
- How much security does the company require?
- What software or hardware does the resource require?
- How much administration does it need?
- How much will it cost?
- Will this resource meet the needs of the organization today and in the future?
- Will additional training be needed?

Networks cannot just be put together at the drop of a hat. A lot of planning is required before implementation of a network to ensure that whatever design is chosen will be effective and efficient, and not just for today but for the future as well. The forethought of the designer will lead to the best network with the least amount of administrative overhead. In each network, it is important that a plan be developed to answer the previous questions. The answers will help the designer choose the type of resource model to use.

Peer-to-Peer Networks

In a peer-to-peer network, the computers act as both service providers and service requestors. An example of a peer-to-peer resource model is shown in Figure 6.7.

The peer-to-peer model is great for small, simple, inexpensive networks. This model can be set up almost instantly, with little extra hardware required. Many versions of Windows (Windows 11, Windows 10, and others) as well as Linux and macOS are popular operating

system environments that support the peer-to-peer resource model. Peer-to-peer networks are also referred to as *workgroups*.

FIGURE 6.7 The peer-to-peer resource model

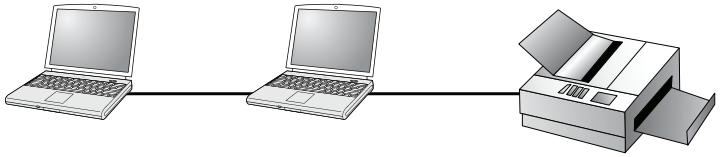

Generally speaking, there is no centralized administration or control in the peer-to-peer resource model. Every workstation has unique control over the resources that the computer owns, and each workstation must be administered separately. However, this very lack of centralized control can make administering the network difficult; for the same reason, the network isn't very secure. Each user needs to manage separate passwords for each computer on which they wish to access resources, as well as set up and manage the shared resources on their own computer. Moreover, because each computer is acting as both a workstation and a server, it may not be easy to locate resources. The person who is in charge of a file may have moved it without anyone's knowledge. Also, the users who work under this arrangement need more training because they are not only users but also administrators.

Will this type of network meet the needs of the organization today and in the future? Peer-to-peer resource models are generally considered the right choice for small companies that don't expect future growth. Small companies that expect growth, on the other hand, should not choose this type of model.

A rule of thumb is that if you have no more than 10–20 computers and centralized security is not a key priority, a workgroup may be a good choice for you.

Client-Server Resource Model

The client-server model (also known as *server-based model*) is better than the peer-to-peer model for large networks (say, more than 10 computers) that need a more secure environment and centralized control. Server-based networks use one or more dedicated, centralized servers. All administrative functions and resource sharing are performed from this point. This makes it easier to share resources, perform backups, and support an almost unlimited number of users.

This model also offers better security than the peer-to-peer model. However, the server needs more hardware than a typical workstation/server computer in a peer-to-peer resource model. In addition, it requires specialized software (the server OS) to manage the server's role in the environment of handling multiple client requests simultaneously. With the addition of a server hardware and OS, server-based networks can easily cost more than

peer-to-peer resource models. However, for large networks, it's the only choice. An example of a client-server resource model is shown in Figure 6.8.

FIGURE 6.8 The client-server resource model

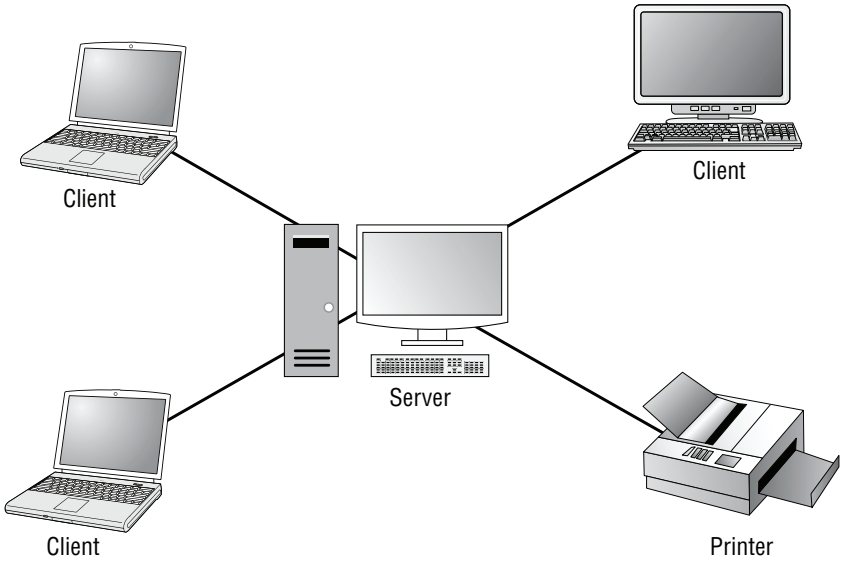

Server-based networks are often known as *domains*. The key characteristic of a server-based network is that security is centrally administered. When you log into the network, the login request is passed to the server responsible for security, sometimes known as a *domain controller*. (Microsoft uses the term domain controller, whereas other vendors of server products do not.) This is different from the peer-to-peer model, where each individual workstation validates users. In a peer-to-peer model, if the user `jsmith` wants to be able to log into different workstations, they need to have a user account set up on each machine. This can quickly become an administrative nightmare! In a domain, all user accounts are stored on the server. User `jsmith` needs only one account and can log into any of the workstations in the domain.

A Microsoft domain is a network of Windows computers that is controlled by a Windows Server running Active Directory—that server is called a domain controller.

Client-server resource models are the desired models for companies that are continually growing, that need to support a large environment, or that need centralized security. Server-based networks offer the flexibility to add more resources and clients almost indefinitely into the future. Hardware costs may be higher, but with the centralized administration, managing resources becomes less time-consuming. Also, only a few administrators need to be trained, and users are responsible for only their own work environment.

If you are looking for an inexpensive, simple network with little setup required, and there is no need for the company to grow in the future, then the peer-to-peer network is the way to go. If you are looking for a network to support many users (more than 10–20 computers), strong security, and centralized administration, consider server-based or cloud-based networks your best choices.

Whatever you decide, always take the time to plan your network before installing it. A network is not something you can just throw together. You don't want to find out a few months down the road that the type of network you chose does not meet the needs of the company—this could be a time-consuming and costly mistake.

Network Topologies

A *topology* is a way of physically laying out the network. When you plan and install a network, you need to choose the right topology for your situation. Each type differs from the others by its cost, ease of installation, fault tolerance (how the topology handles problems such as cable breaks), and ease of reconfiguration (such as adding a new workstation to the existing network).

There are five primary topologies:

- Bus
- Star (or hub-and-spoke)
- Ring
- Mesh
- Hybrid

Each topology has advantages and disadvantages. Table 6.1 summarizes the advantages and disadvantages of each topology, and then we will go into more detail about each one.

TABLE 6.1 Topologies—advantages and disadvantages

Topology	Advantages	Disadvantages
Bus	Cheap. Easy to install.	Difficult to reconfigure. A break in the bus disables the entire network.
Star	Cheap. Very easy to install and reconfigure. More resilient to a single cable failure.	More expensive than bus.
Ring	Efficient. Easy to install.	Reconfiguration is difficult. Very expensive.
Mesh	Best fault tolerance.	Reconfiguration is extremely difficult, extremely expensive, and very complex.
Hybrid	Gives a combination of the best features of each topology used.	Complex (less so than mesh, however).

Bus Topology

A *bus topology* is the simplest. It consists of a single cable that runs to every workstation, as shown in Figure 6.9. This topology uses the least amount of cabling. Each computer shares the same data and address path. With a bus topology, messages pass through the trunk, and each workstation checks to see if a message is addressed to it. If the address of the message matches the workstation's address, the network adapter retrieves it. If not, the message is ignored.

FIGURE 6.9 The bus topology

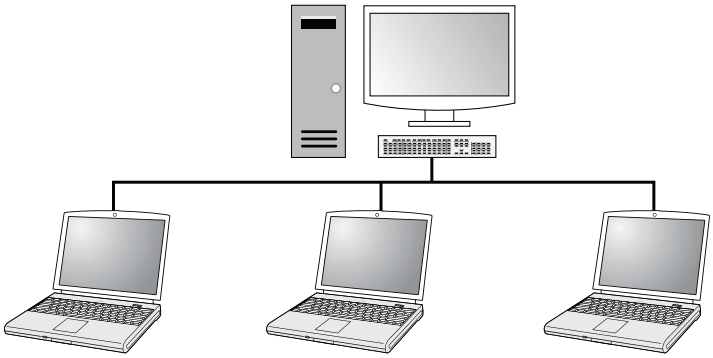

Cable systems that use the bus topology are easy to install. You run a cable from the first computer to the last computer. All of the remaining computers attach to the cable somewhere in between. Because of the simplicity of installation, and because of the low cost of the cable, bus topology cabling systems are the cheapest to install.

Although the bus topology uses the least amount of cabling, it is difficult to add a workstation. If you want to add another workstation, you have to reroute the cable completely and possibly run two additional lengths of it. Also, if any one of the cables breaks, the entire network is disrupted. Therefore, such a system is expensive to maintain and can be difficult to troubleshoot. You will rarely run across physical bus networks in use today.

Star Topology

A *star topology* (also called a *hub-and-spoke topology*) branches each network device off a central device called a *hub* or a *switch*, making it easy to add a new workstation. If a workstation goes down, it does not affect the entire network; if the central device goes down, the entire network goes with it. Because of this, the hub (or switch) is called a *single point of failure*. Figure 6.10 shows a simple star network.

Star topologies are very easy to install. A cable is run from each workstation to the switch. The switch is placed in a central location in the office (e.g., a utility closet). Star topologies are more expensive to install than bus networks because several more cables need to be installed, plus the switches. But the ease of reconfiguration and fault tolerance (one cable failing does not bring down the entire network) far outweigh the drawbacks. This is by far the most commonly installed network topology in use today.

FIGURE 6.10 The star topology

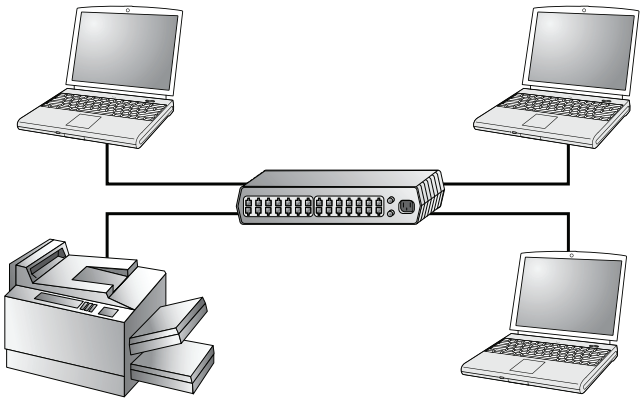

Although the switch is the central portion of a star topology, some older networks use a device known as a hub instead of a switch. Switches are more advanced than hubs, and they provide better performance than hubs for only a small price increase. Colloquially, though, many administrators use the terms *hub* and *switch* interchangeably.

Ring Topology

In a *ring topology*, each computer connects to two other computers, joining them in a circle and creating a unidirectional path where messages move from workstation to workstation. Each entity participating in the ring reads a message and then regenerates it and hands it to its neighbor on a different network cable. See Figure 6.11 for an example of a ring topology.

FIGURE 6.11 The ring topology

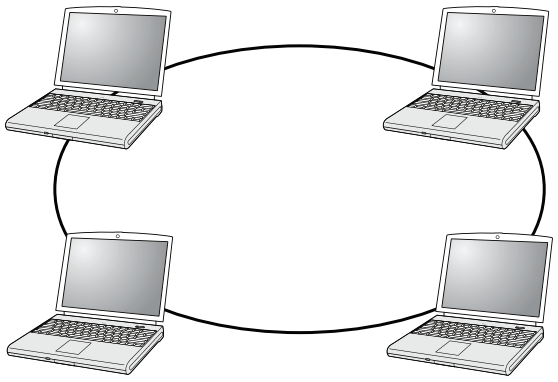

The ring makes it difficult to add new computers. Unlike a star topology network, a ring topology network will go down if one entity is removed from the ring. Physical ring

topology systems rarely exist anymore, mainly because the hardware involved was fairly expensive and the fault tolerance was very low.

You might have heard of an older network architecture called Token Ring. Contrary to its name, it does *not* use a physical ring. It actually uses a physical star topology, but the traffic flows in a logical ring from one computer to the next.

Mesh Topology

The *mesh topology* is the most complex in terms of physical design. In this topology, each device is connected to every other device (see Figure 6.12). This topology is rarely found in wired LANs, mainly because of the complexity of the cabling. If there are x computers, there will be $(x \times (x - 1)) \div 2$ cables in the network. For example, if you have five computers in a mesh network, it will use $(5 \times (5 - 1)) \div 2 = 10$ cables. This complexity is compounded when you add another workstation. For example, your 5-computer, 10-cable network will jump to 15 cables if you add just one more computer. Imagine how the person doing the cabling would feel if you told them they had to cable 50 computers in a mesh network—they'd have to come up with $(50 \times (50 - 1)) \div 2 = 1{,}225$ cables! (Not to mention figuring out how to connect them all.)

FIGURE 6.12 The mesh topology

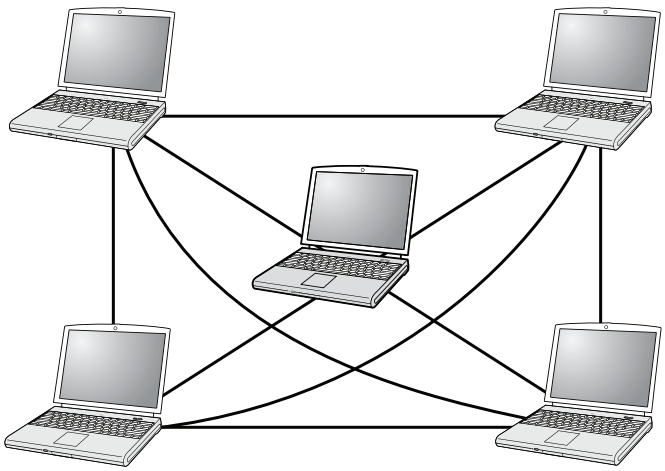

Because of its design, the physical mesh topology is expensive to install and maintain. Cables must be run from each device to every other device. The advantage you gain is high fault tolerance. With a mesh topology, there will always be a way to get the data from source to destination. The data may not be able to take the direct route, but it can take an alternate, indirect route. For this reason, the mesh topology is often used to

connect multiple sites across WAN links. It uses devices called *routers* to search multiple routes through the mesh and determine the best path. However, the mesh topology does become inefficient with five or more entities because of the number of connections that need to be maintained.

Hybrid Topology

The *hybrid topology* is simply a mix of the other topologies. It would be impossible to illustrate it because there are many combinations. In fact, most networks today are not only hybrid but heterogeneous. (They include a mix of components of different types and brands.) The hybrid network may be more expensive than some types of network topologies, but it takes the best features of all the other topologies and exploits them.

Table 6.1, earlier in this chapter, summarizes the advantages and disadvantages of each type of network topology.

Rules of Communication

Regardless of the type of network you choose to implement, the computers on that network need to know how to talk to each other. To facilitate communication across a network, computers use a common language called a *protocol*. We'll cover protocols more in Volume 1 Chapter 7, "Introduction to TCP/IP," but for now know that essentially they are languages much like English is a language. Within each language, there are rules that need to be followed so that all computers understand the right communication behavior.

To use a human example, within English there are grammar rules. If you put a bunch of English words together in a way that doesn't make sense, no one will understand you. If you just decide to omit verbs from your language, you're going to be challenged to get your point across. And if everyone talks at the same time, the conversation can be hard to follow.

Computers need standards to follow to keep their communication clear. Different standards are used to describe the rules that computers need to follow to communicate with each other. The most important communication framework, and the backbone of all networking, is the OSI model.

The OSI model is not specifically listed in the CompTIA A+ exam objectives. However, it's a critical piece of networking knowledge and a framework with which all technicians should be familiar. For example, you might hear (on the exam or in real life) someone talking about a "Layer 3" device, and you need to understand the implications of that!

OSI Model

The International Organization for Standardization (ISO) published the *Open Systems Interconnection (OSI)* model in 1984 to provide a common way of describing network protocols.

The ISO put together a seven-layer model providing a relationship between the stages of communication, with each layer adding to the layer above or below it.

This OSI model is a theoretical model governing computer communication. Even though at one point an "OSI protocol" was developed, it never gained wide acceptance. You will never find a network that is running the "OSI protocol."

Here's how the theory behind the OSI model works. As a transmission takes place, the higher layers pass data through the lower layers. As the data passes through a layer, that layer tacks its information (also called a *header*) onto the beginning of the information being transmitted until it reaches the bottom layer. A layer may also add a trailer to the end of the data. The bottom layer sends the information out on the wire (or in the air, in the case of wireless).

At the receiving end, the bottom layer receives and reads the information in the header, removes the header and any associated trailer related to its layer, and then passes the remainder to the next highest layer. This procedure continues until the topmost layer receives the data that the sending computer sent.

The OSI model layers are listed here from top to bottom, with descriptions of what each of the layers is responsible for:

7—Application Layer The Application layer allows access to network services. This is the layer at which file services, print services, and other applications operate.

6—Presentation Layer This layer determines the "look," or format, of the data. The Presentation layer performs protocol conversion and manages data compression, data translation, and encryption. The character set information also is determined at this level. (The *character set* determines which numbers represent which alphanumeric characters.)

5—Session Layer This layer allows applications on different computers to establish, maintain, and end a session. A *session* is one virtual conversation. For example, all the procedures needed to transfer a single file make up one session. Once the session is over, a new process begins. This layer enables network procedures, such as identifying passwords, logins, and network monitoring.

4—Transport Layer The Transport layer controls the data flow and troubleshoots any problems with transmitting or receiving datagrams. It also takes large messages and segments them into smaller ones and takes smaller segments and combines them into a single, larger message, depending on which way the traffic is flowing. Finally, the TCP protocol (one of the two options at this layer) has the important job of verifying that the destination host has received all packets, providing error checking and reliable end-to-end communications.

3—Network Layer The Network layer is responsible for logical addressing of messages. At this layer, the data is organized into chunks called *packets*. The Network layer

is something like the traffic cop. It is able to judge the best network path for the data based on network conditions, priority, and other variables. This layer manages traffic through packet switching, routing, and controlling congestion of data.

2—Data Link Layer This layer arranges data into chunks called *frames*. Included in these chunks is control information indicating the beginning and end of the data stream. The Data Link layer is very important because it makes transmission easier and more manageable, and it allows for error checking within the data frames. The Data Link layer also describes the unique physical address (also known as the *MAC address*) for each NIC. The Data Link layer is subdivided into two sections: Media Access Control (MAC) and Logical Link Control (LLC).

1—Physical Layer The Physical layer describes how the data gets transmitted over a communication medium. This layer defines how long each piece of data is and the translation of each into the electrical pulses or light impulses that are sent over the wires, or the radio waves that are sent through the air. It decides whether data travels unidirectionally or bidirectionally across the hardware. It also relates electrical, optical, mechanical, and functional interfaces to the cable.

Figure 6.13 shows the complete OSI model. Note the relationship of each layer to the others and the function of each layer.

A helpful mnemonic device to remember the OSI layers in order is "All People Seem To Need Data Processing."

IEEE 802 Standards

Continuing with our theme of communication, it's time to introduce one final group of standards. You've already learned that a protocol is like a language; think of the IEEE 802 standards as syntax, or the rules that govern who communicates, when they do it, and how they do it.

The IEEE formed a subcommittee to create standards for network types. These standards specify certain types of networks, although not every network protocol is covered by the IEEE 802 committee specifications. This model contains several standards. The ones commonly in use today are 802.3 CSMA/CD (Ethernet) LAN and 802.11 Wireless networks. The IEEE 802 standards were designed primarily for enhancements to the bottom three layers of the OSI model. The IEEE 802 standard breaks the Data Link layer into two sublayers: a Logical Link Control (LLC) sublayer and a Media Access Control (MAC) sublayer. The Logical Link Control sublayer manages data link communications. The Media Access Control sublayer watches out for data collisions and manages physical addresses, also referred to as MAC addresses.

You've most likely heard of 802.11be (Wi-Fi 7), 802.11ax (Wi-Fi 6), 802.11ac (Wi-Fi 5), or 802.11n wireless networking. The rules for communicating with all versions of 802.11 are defined by the IEEE standard. Another very well-known standard is 802.3 CSMA/CD. You might know it by its more popular name, Ethernet.

FIGURE 6.13 The OSI model

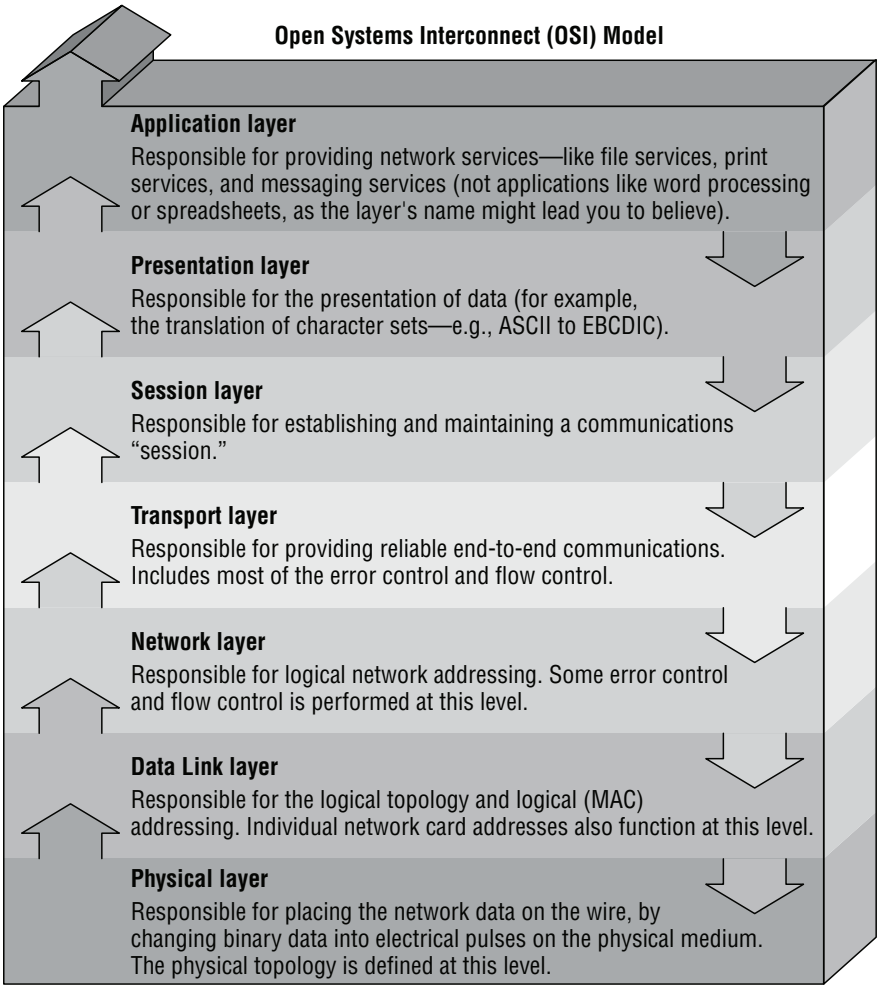
Open Systems Interconnect (OSI) Model
Application layer
Responsible for providing network services—like file services, print services, and messaging services (not applications like word processing or spreadsheets, as the layer's name might lead you to believe).
Presentation layer
Responsible for the presentation of data (for example, the translation of character sets—e.g., ASCII to EBCDIC).
Session layer
Responsible for establishing and maintaining a communications "session."
Transport layer
Responsible for providing reliable end-to-end communications. Includes most of the error control and flow control.
Network layer
Responsible for logical network addressing. Some error control and flow control is performed at this level.
Data Link layer
Responsible for the logical topology and logical (MAC) addressing. Individual network card addresses also function at this level.
Physical layer
Responsible for placing the network data on the wire, by changing binary data into electrical pulses on the physical medium. The physical topology is defined at this level.

The original 802.3 CSMA/CD standard defines a bus topology network that uses a 50-ohm coaxial baseband cable and carries transmissions at 10 Mbps. This standard groups data bits into frames and uses the *Carrier Sense Multiple Access with Collision Detection (CSMA/CD)* cable access method to put data on the cable. Currently, the 802.3 standard has been amended to include speeds up to 400 Gbps over multimode fiber-optic cable.

Breaking the CSMA/CD acronym apart may help illustrate how it works:

CS First, there is the Carrier Sense (CS) part, which means that computers on the network are listening to the wire at all times.

MA Multiple Access (MA) means that multiple computers have access to the line at the same time. This is analogous to having five people on a conference call. Everyone is listening, and everyone in theory can try to talk at the same time. Of course, when more

than one person talks at once, there is a communication error. In CSMA/CD, when two machines transmit at the same time, a data *collision* takes place and the intended recipients receive none of the data.

CD This is where the Collision Detection (CD) portion of the acronym comes in; the collision is detected and each sender knows they need to send again. Each sender then waits for a short, random period of time and tries to transmit again. This process repeats until transmission takes place successfully.

The CSMA/CD technology is considered a *contention-based* access method.

The only major downside to 802.3 is that with large networks (more than 100 computers on the same segment), the number of collisions increases to the point where more collisions than transmissions are taking place.

Other examples of contention methods exist, such as *Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA)*. Whereas CSMA/CD tries to fix collisions after they happen, CSMA/CA tries to avoid them in the first place by actively listening and only transmitting when the channel is clear. Wireless Ethernet uses CSMA/CA.

The alternative to the contention-based method is a controlled access method. In controlled access, the only computer that can transmit is the one given permission, such as by being granted a token in the aforementioned Token Ring network. Once it's done transmitting, it passes the token to the next computer, which can then transmit.

Identifying Common Network Hardware

We have looked at the types of networks, network topologies, and the way communications are handled. That's all of the logical stuff. To really get computers to talk to each other requires hardware. Every computer on the network needs to have a network adapter of some type. In many cases, you also need some sort of cable to hook them together. (Wireless networking is the exception, but at the backend of a wireless network there are still components wired together.) And finally, you might also need connectivity devices to attach several computers or networks to each other. We'll look at all of these in this section, starting with the component closest in to the "local computer," or the one you're at, and working outward.

Network Interface Cards

You were introduced to the *network interface card (NIC)*, also referred to as a network adapter card, earlier in the chapter. It provides the physical interface between computer and

cabling, and prepares data, sends data, and controls the flow of data. It can also receive and translate data into bytes for the CPU to understand. NICs come in many shapes and sizes.

Different NICs are distinguished by the PC bus type and the network for which they are used. The following sections describe the role of NICs and how to evaluate them.

Exam Objective 2.5, "Compare and contrast common networking hardware devices," lists 10 different networking devices to know. The first is the NIC.

Compatibility

The first thing you need to determine is whether the NIC will fit the bus type of your PC. If you have more than one type of bus in your PC (for example, a combination PCI/PCI Express), use a NIC that fits into the fastest type (the PCI Express, in this case). This is especially important in servers because the NIC can quickly become a bottleneck if this guideline isn't followed.

More and more computers are using NICs that have USB interfaces. For the rare laptop computer that doesn't otherwise have a NIC built into it, these small portable cards are very handy.

A USB network card can also be handy for troubleshooting. If a laptop isn't connecting to the network properly with its built-in card, you may be able to use the USB NIC to see if it's an issue with the card or perhaps a software problem.

Network Interface Card Performance

The most important goal of the NIC is to optimize network performance and minimize the amount of time needed to transfer data packets across the network. The key is to ensure that you get the fastest card that you can for the type of network that you're on. For example, if your wireless network supports 802.11ac/ax/be, make sure to get an 802.11be card because it's the fastest.

Sending and Controlling Data

In order for two computers to send and receive data, the cards must agree on several things:

- The maximum size of the data frames
- The amount of data sent before giving confirmation
- The time needed between transmissions
- The amount of time to wait before sending confirmation
- The speed at which data transmits

If the cards can agree, the data is sent successfully. If the cards cannot agree, the data is not sent.

To send data on the network successfully, all NICs need to use the same media access method (such as CSMA/CD) and be connected to the same piece of cable. This usually isn't a problem, because the vast majority of network cards sold today are Ethernet.

In addition, NICs can send data using either full-duplex or half-duplex mode. *Half-duplex communication* means that between the sender and receiver, only one of them can transmit at any one time. In *full-duplex communication*, a computer can send and receive data simultaneously. The main advantage of full-duplex over half-duplex communication is performance. NICs (Gigabit Ethernet NICs) can operate twice as fast (1 Gbps) in full-duplex mode as they do normally in half-duplex mode (500 Mbps). In addition, collisions are avoided, which speeds up performance as well. Configuring the network adapter's duplexing setting is done from the Advanced tab of the NIC's properties, as shown in Figure 6.14.

Normally, you aren't going to have to worry about how your NIC sends or controls data. Just make sure to get the fastest NIC that is compatible with your network. Do know that the negotiations discussed here are happening in the background, though.

FIGURE 6.14 A NIC's Speed & Duplex setting

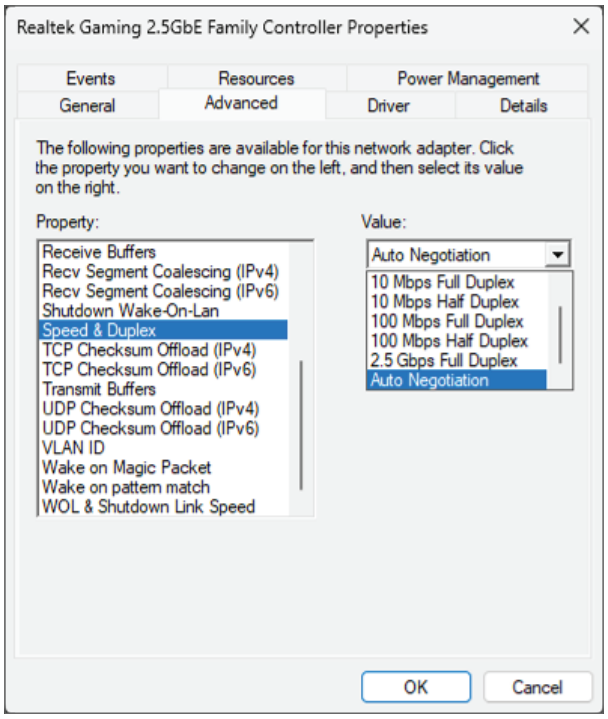

NIC Configuration

Each card must have a unique hardware address, called a *Media Access Control address* or *MAC address*. (Remember earlier in the chapter we said you didn't need to know the OSI model for the exam, but you should know it anyway? Here's an example of why. Now you can piece together that this is a physical address that is referenced at Layer 2.) If two NICs on the same network have the same hardware address, neither one will be able to communicate. For this reason, the IEEE has established a standard for hardware addresses and assigns blocks of these addresses to NIC manufacturers, who then hardwire the addresses into the cards.

Know what the MAC address is for the exam. It's part of Exam Objective 2.5.

MAC addresses are 48 bits long and written in hexadecimal, such as `40-61-86-E4-5A-9A`. An example is shown in Figure 6.15 from the output of the `ipconfig /all` command executed at the Windows command prompt. In macOS, the analogous command is `ifconfig` (the MAC address may be labeled `HWaddr`). In Linux, the command is simply `ip`.

Although it is possible for NIC manufacturers to produce multiple NICs with the same MAC address, it happens very rarely. If you do encounter this type of problem, contact the hardware manufacturer.

FIGURE 6.15 Physical (MAC) address

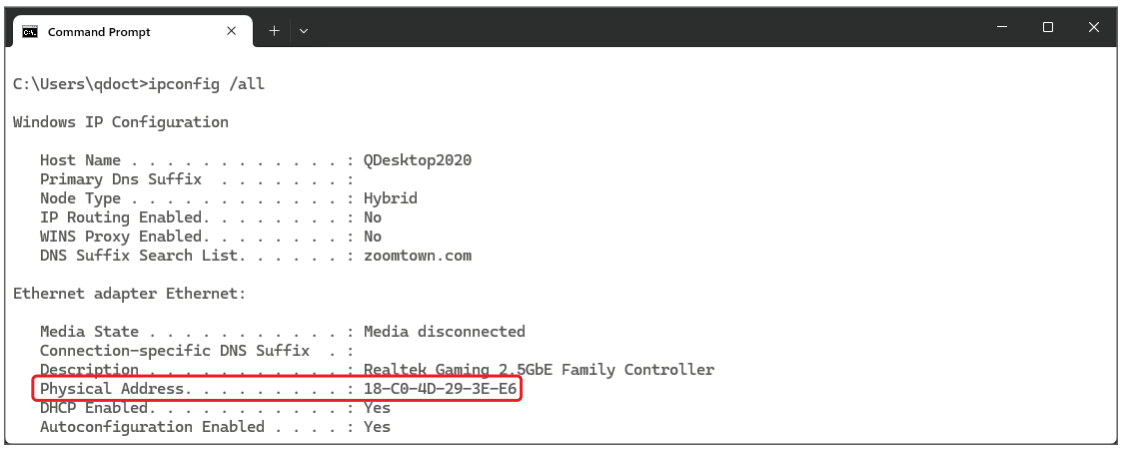
Command Prompt
C:\Users\qdoct>ipconfig /all
Windows IP Configuration
Host Name . . . . . : QDesktop2020
Primary Dns Suffix . . . . . :
Node Type . . . . . : Hybrid
IP Routing Enabled. . . . . : No
WINS Proxy Enabled. . . . . : No
DNS Suffix Search List. . . . . : zoomtown.com
Ethernet adapter Ethernet:
Media State . . . . . : Media disconnected
Connection-specific DNS Suffix . :
Description . . . . . : Realtek Gaming 2.5GbE Family Controller
Physical Address. . . . . : 18-C0-4D-29-3E-E6
DHCP Enabled. . . . . : Yes
Autoconfiguration Enabled . . . . : Yes

Think of the MAC address as being physically burned into the network card and that it can't be changed, unlike logical IP network addresses, which are easily modified. (We will cover IP addresses in depth in Volume 1 Chapter 7.) We say this even though it's sort of not technically accurate. It's true that MAC addresses are permanently assigned to NICs.

However, MAC addresses *can* be changed by software controlling the NIC—it's called *spoofing*. Hackers can use spoofing to impersonate a different NIC for anonymity or gaining unauthorized access to a network. There may be legitimate reasons to spoof a MAC as well. For example, say you have a networking software package tied to a specific server's NIC, and the NIC dies. To keep the software package running, you can install a new NIC and spoof the old one's MAC address, until you have time to contact the software manufacturer to resolve the issue. All of this said, for the A+ exam it's still best to think of a MAC address as a permanent NIC identifier that can't be changed.

NIC Drivers

In order for the computer to use the NIC, it is very important to install the proper device drivers. Recall that you learned in Volume 1 Chapter 2, "Expansion Cards, Storage Devices, and Power Supplies," that drivers are pieces of software that communicate directly with the operating system. Drivers are specific to each NIC and operating system, and they operate in the Media Access Control (MAC) sublayer of the Data Link layer of the OSI model.

To see which version the driver is, you need to look at the device's properties. There are several ways to do this. A common one is to open Device Manager (type **Device Manager** into the Windows search box and press Enter), and find the device in the Network adapters section, as shown in Figure 6.16.

FIGURE 6.16 Device Manager

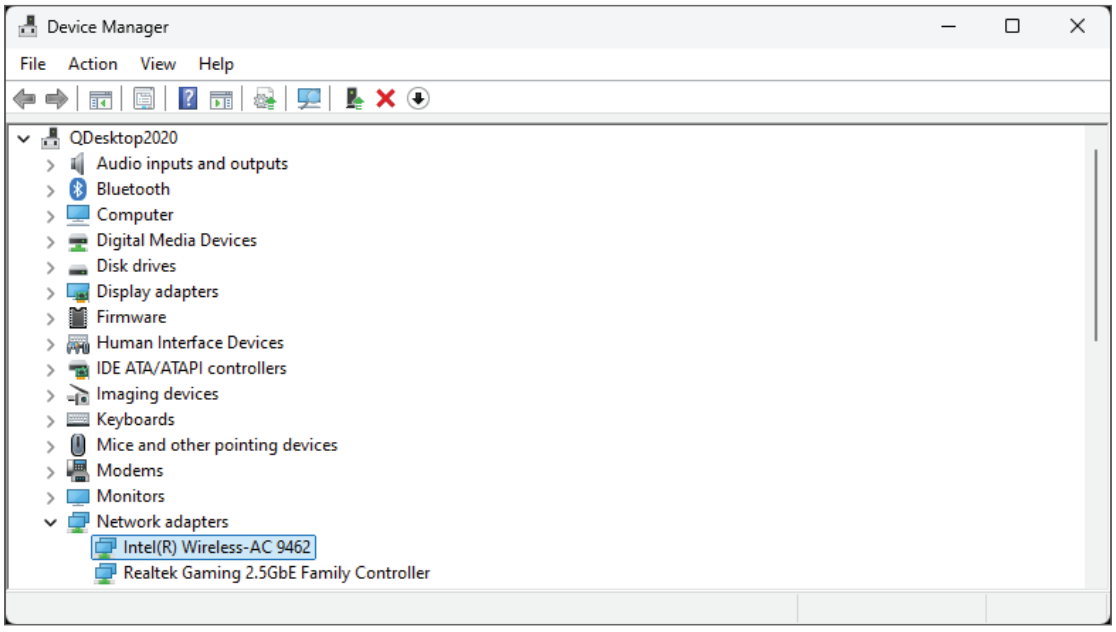

Right-click the device, click Properties, and then go to the Driver tab, as shown in Figure 6.17. Here you can see a lot of information about the driver, update it, or roll

it back if you installed a new one and it fails for some reason. You can also update the driver by right-clicking the device in Device Manager and choosing Update Driver from the menu.

The best place to get drivers is almost always the manufacturer's website. When you click Update Driver, Windows will ask you if you want to search for the driver on the Internet or provide a location for it. The best course of action is to download the driver first, and then tell Windows where you put it.

FIGURE 6.17 NIC properties Driver tab

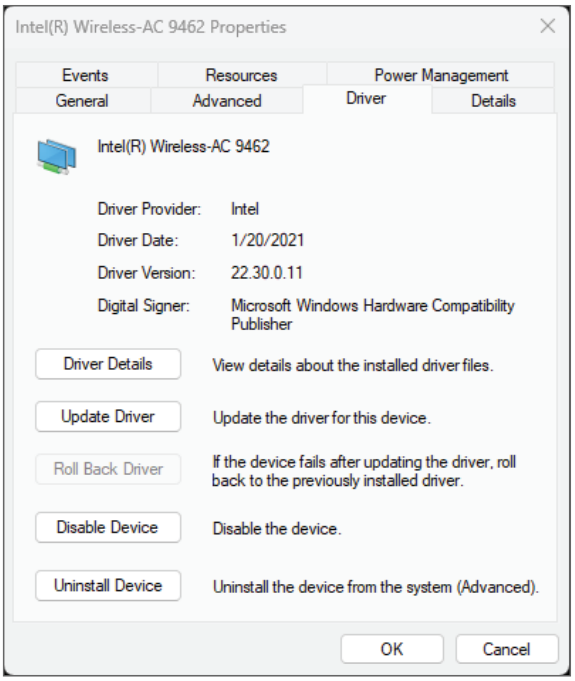

Cables and Connectors

When the data is passing through the OSI model and reaches the Physical layer, it must find its way onto the medium that is used to transfer data physically from computer to computer. This medium is called the *cable* (or in the case of wireless networks, the air). It is the NIC's role to prepare the data for transmission, but it is the cable's role to move the data properly to its intended destination. The following sections discuss the three main types of physical cabling: coaxial, twisted pair, and fiber-optic. (Wireless communication is covered in Volume 1 Chapter 9.)

Coaxial Cable

Coaxial cable (or coax) contains a center conductor core made of copper, which is surrounded by a plastic jacket with a braided shield over it (as shown in Figure 6.18). Either Teflon or a plastic outer coating covers this metal shield.

Exam Objective 3.2, "Summarize basic cable types and their connectors, features, and purposes," lists three subobjectives covered in this section. They are coaxial cable, plenum-rated cable, and F-type connector.

FIGURE 6.18 Coaxial cable

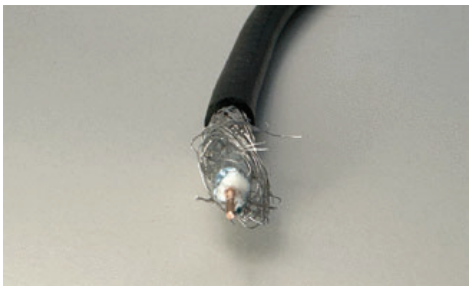

Common network cables are covered with a plastic called *polyvinyl chloride (PVC)*. Although PVC is flexible, fairly durable, and inexpensive, it has a nasty side effect in that it produces poisonous gas when burned. An alternative is a Teflon-type covering that is frequently referred to as a *plenum-rated* coating. Teflon coating does not produce toxic gas when burned and is rated for use in the ventilation areas of a building, called the plenum, that circulate breathable air. Plenum-rated cable is more expensive, but it may be mandated by electrical code whenever cable is hidden in walls or ceilings.

Plenum rating can apply to all types of network cabling.

Coax Cable Specifications

Coaxial cable is available in various specifications that are rated according to the Radio Guide (RG) system, which was originally developed by the U.S. military. The thicker the copper, the farther a signal can travel—and with that comes a higher cost and a less flexible cable. Coax is rarely seen in computer networking today because it's painfully slow; its heyday was several decades ago.

When coax cable was popular for networking, there were two standards that had high use: RG-8 (thicknet) and RG-58A/U (thinnet). Thicknet had a maximum segment distance of 500 meters and was used primarily for network backbones. Thinnet was more often used in a conventional physical bus. A thinnet segment could span 185 meters. Both thicknet and thinnet had impedance of 50 ohms. Table 6.2 shows the different types of RG cabling and

their uses. Although coax is an A+ exam objective, no specific coax cabling standards are currently specified on the exam objectives. The ones that used to be named on the A+ exam objectives were RG-6 and RG-59.

Explaining Ethernet Naming Standards

In Table 6.2, you will notice two terms that might be new to you: *10Base5* and *10Base2*. These are Ethernet naming standards. The number at the beginning tells you the maximum speed that the standard supports, which is 10 Mbps in this case. The word *Base* refers to the type of transmission, either baseband (one signal at a time per cable) or broadband (multiple signals at the same time on one cable). Legend has it that the 5 and the 2 refer to the approximate maximum transmission distance (in hundreds of meters) for each specification. Another usage you will see is *10BaseT*, which refers to twisted pair cabling.

TABLE 6.2 Coax RG types

RG #	Popular name	Ethernet implementation	Max distance	Core type
RG-6	Satellite/cable TV, cable modems	N/A	304 m	Solid copper
RG-8	Thicknet	10Base5	500 m	Solid copper
RG-58 U	N/A	None	185 m	Solid copper
RG-58 A/U	Thinnet	10Base2	185 m	Stranded copper
RG-59	Cable television	N/A	228 m	Solid copper

Coaxial networking has all but gone the way of the dinosaur. The only two coaxial cable types you might see today are RG-6 and RG-59, which are used with cable television. Of the two, RG-6 has a thicker core (1.0 mm), can run longer distances (up to 304 meters, or 1,000 feet), and can support digital signals. RG-59 (0.762 mm core) is considered adequate for analog cable TV but not digital and has a maximum distance of about 228 meters (750 feet). The maximum speed for each depends on the quality of the cable and the standard on which it's being used. Both have impedance of 75 ohms.

Coax Connector Types

Thicknet was a bear to use. Not only was it highly inflexible, but you also needed to use a connector called a *vampire tap*. A vampire tap is so named because a metal tooth sinks into the cable, thus making the connection with the inner conductor. The tap is connected to an external transceiver that in turn has a 15-pin AUI connector (also called a *DIX* or *DB-15* connector) to which you attach a cable that connects to the station. The transceiver is shown in Figure 6.19. On the right side, you will see the thicknet cable running through the portion of the unit that contains the vampire tap. DIX got its name from the companies that worked on this format—Digital, Intel, and Xerox.

FIGURE 6.19 Thicknet transceiver and cable inside a vampire tap

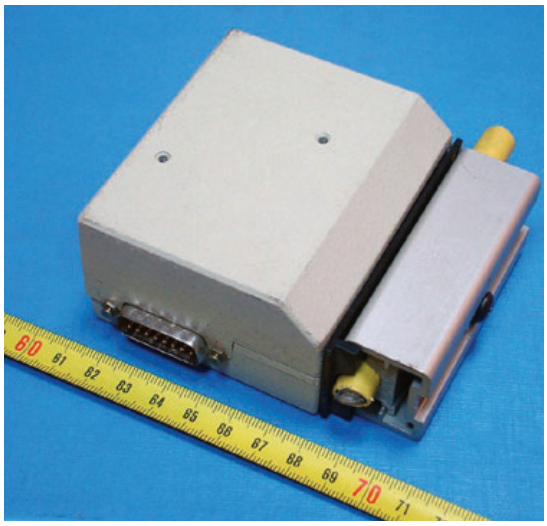

Thicknet transceiver licensed Under CC BY-Sa 2.5 via Wikimedia Commons. http://commons.wikimedia.org/wiki/File:ThicknetTransceiver.jpg#/media/File:ThicknetTransceiver.jpg.

Thinnet coax was much easier to use. Generally, thinnet networking cables used a *BNC connector* (see Figure 6.20) to attach to a T-shaped connector that attached to the workstation. The other side of the T-connector would either continue on with another thinnet segment or be capped off with a terminator. Figure 6.21 shows what it looks like all connected—the open male connector pointing towards the camera would connect to the NIC. It is beyond the scope of this book to settle the long-standing argument over the meaning of the abbreviation BNC. We have heard Bayonet Connector, Bayonet Nut Connector, and British Naval Connector—among others. What is relevant is that the BNC connector locks securely with a quarter-twist motion.

Another type of connector that you will see in use with coax is a *splitter*. As its name implies, a splitter takes a single signal (say that three times fast) and splits it into multiple replicas of the same signal. You might use this for cable TV—one line may run into your

house, but the signal ultimately needs to get split for three televisions. This type of configuration will work for cable TV or cable Internet. Figure 6.22 shows a one-to-two coax splitter. You can also buy splitters that split one input into three or more outputs.

FIGURE 6.20 Two T-connectors, male BNC connector, and two terminators

FIGURE 6.21 Attached coax cables and connectors

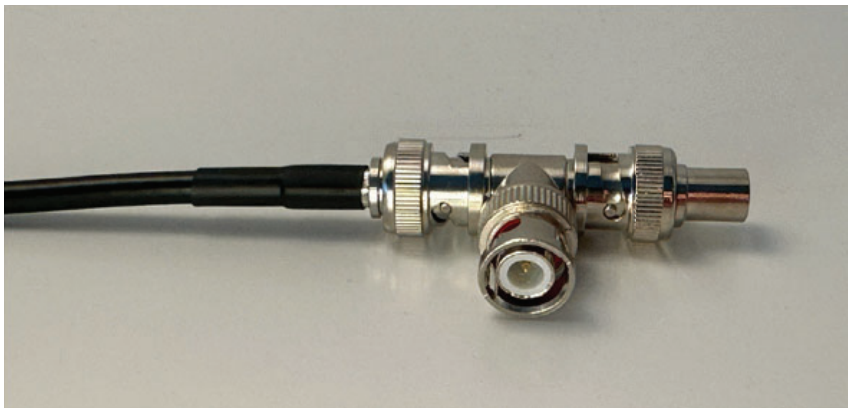

Keep in mind that a coax signal is designed to go from one sender to one receiver, so splitting it can cause some issues. Splitting the signal causes it to weaken, meaning that signal quality could be lower, and it might not travel the same distance as a non-split signal. To avoid problems, don't over-split the cable, and purchase a good-quality or amplified splitter.

The last type of coax connector we will cover is called an *F-connector* (referred to in exam objectives as an *F-type connector*, shown in Figure 6.23 to the left of a BNC connector), and it is used with cable TV. You'll see it on the end of an RG-6 or possibly an RG-59 cable. The exposed end of the copper cable is pushed into the receptacle, and the connector is threaded so that it can screw into place.

FIGURE 6.22 A coax splitter

FIGURE 6.23 F-type and BNC connectors

Twisted Pair Cable

Twisted pair is the most popular type of cabling to use because of its flexibility and low cost. It consists of several pairs of copper wire twisted around each other within an insulated jacket, as shown in Figure 6.24.

Exam Objective 3.2 has several subobjectives related to twisted pair cable that are covered in this section. Be familiar with shielded twisted pair and unshielded twisted pair, direct burial, cable categories, T568A/T568B standards, and RJ11 and RJ45 connectors.

There are two different types of twisted pair cables: *shielded twisted pair (STP)* and *unshielded twisted pair (UTP)*. Both types of cable have two or four pairs of twisted wires going through them. The difference is that STP has an extra layer of braided foil shielding surrounding the wires to decrease electrical interference, as shown in Figure 6.25.

(In Figure 6.25, the individual wire pairs are shielded as well.) UTP has a PVC or plenum coating but no outer foil shield to protect it from interference. In the real world, UTP is the most common networking cable type used. STP has been used less frequently, but newer standards rely on shielding to deliver ultra-fast transmission speeds.

FIGURE 6.24 Unshielded twisted pair cable

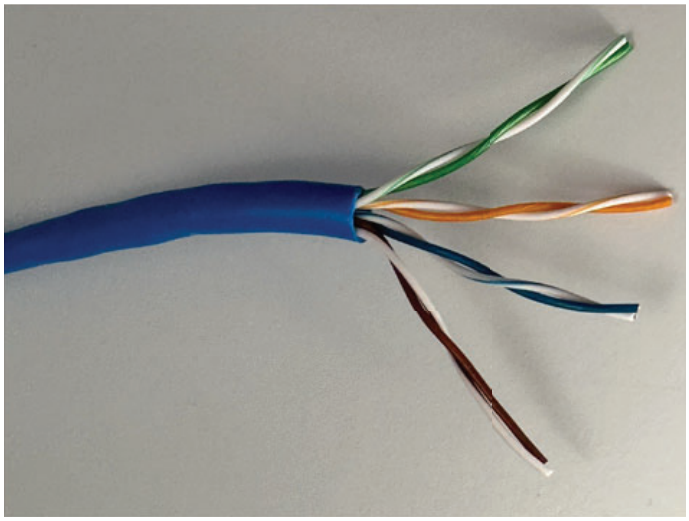

FIGURE 6.25 Shielded twisted pair cable

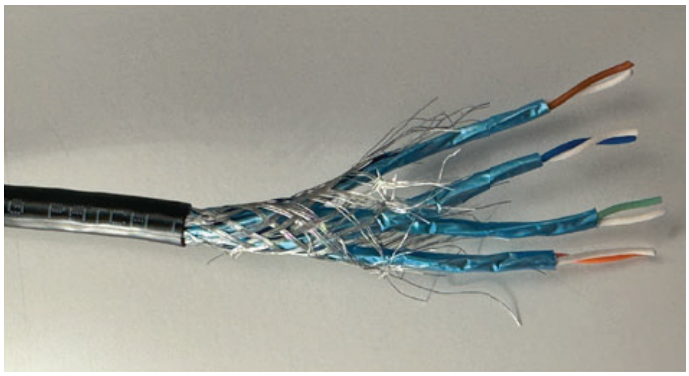

You will often hear people refer to UTP cables as Ethernet cables. As you learned earlier in this chapter, Ethernet is an access method based on the IEEE 802.3 standard and not related to a specific cable type. So while technically it's incorrect, there's no harm in using the term to refer to the cable.

Twisted Pair Cable Specifications

Twisted pair cabling has been in use, at least with old analog telephone lines, for a few generations now. Over time, the need for higher transmission speeds required faster cabling, and the cable manufacturers have been up to the challenge. Now you can find twisted pair in several categories to offer different levels of performance and protection against electrical interference:

Category 5 (Cat 5) came out in 1995 and is able to transmit data at speeds up to 100 Mbps. It was hugely popular.

Category 5e (Cat 5e) was the first copper network cable able to transmit data at speeds up to 1 Gbps. It's basically Cat 5, except that the four twisted pairs of copper wire are physically separated and contain more twists per foot. This provides better interference protection.

Category 6 (Cat 6) is able to transmit data at speeds up to 10 Gbps, but only up to a distance of 55 meters. Its four twisted pairs of copper wire are oriented differently than in Cat 5e. This is the lowest grade of cable you should ever use as a backbone to connect different parts of a network together, such as those on different floors of a building.

Category 6a (Cat 6a) can also handle 10 Gbps speed, but at longer distances (up to 100 meters) than Cat 6 can.

Category 7 (Cat 7) is an incremental upgrade over Cat 6a. Its most notable feature is that every wire pair is shielded, which provides better resistance to crosstalk and external noise. It can handle 10 Gbps at up to 100 meters. You might hear it also called Class F cabling.

Category 8 (Cat 8) improves upon Cat 7 by offering speeds of 25 Gbps or 40 Gbps at up to 30 meters, and it supports 10 Gbps at 100 meters. It's always shielded, just as Cat 7 is. There are two types of Class 8 cable. Class I Cat 8 (also known as Cat 8.1) uses RJ45 connectors and is backward-compatible with all other twisted pair installations using RJ45 connectors. Class II Cat 8 (Cat 8.2) uses different connectors (called twisted pair enhanced registered jack [TERA]) and is not compatible with RJ45 installations.

For as long as twisted pair has existed, every technician has needed to memorize its standard maximum transmission distance of 100 meters (328 feet). You should burn that into your brain, too. Note, however, that some newer standards have shorter maximum distances. For example, if you want to run 10GBaseT over Cat 6, you won't get that much distance—about 55 meters under ideal conditions. Cat 8 (which isn't a specific exam objective) can provide up to 40 Gbps but only at 30 meters.

A Few Notes on Twisted Pair Cabling

CompTIA (and many others) usually shortens the word *category* to Cat and, for example, uses the form Cat 5 to refer to Category 5. This is a common way to refer to these categories, and you can feel free to use these terms interchangeably. If you are buying cable today, you shouldn't buy anything older than Cat 5e.

Before the Cat 6a standard was finalized, several cable vendors started selling Cat 6e cable. Much like Cat 5e was an enhanced version of Cat 5, Cat 6e was supposed to be an enhanced version of Cat 6. The thing is, an official Cat 6e specification was never officially established; if you see Cat 6e, there is no guarantee of the type of performance that you will get. The official standard is Cat 6a.

You'll also see the categories written in different ways. For example, Cat 6a, Cat-6a, or CAT 6a. Different companies use different abbreviation standards but they're all the same thing.

For the exam, you should know at least the specifications for Cat 5, Cat 5e, Cat 6, and Cat 6a. In real life, you should be familiar with basic performance data for all standards Cat 5e and newer. We've summarized the standards in Table 6.3 to make it easier to see them all side by side.

TABLE 6.3 Common twisted pair cabling standards

Standard	Maximum bandwidth	Maximum distance
Cat 3	10 Mbps	100 m
Cat 5	100 Mbps	100 m
Cat 6	10 Gbps	55 m
Cat 6a	10 Gbps	100 m
Cat 7	10 Gbps	100 m
Cat 8	40 Gbps or 10 Gbps	30 m or 100 m

Twisted Pair Connector Types

Twisted pair cabling uses a connector type called an *RJ (registered jack)* connector. Most landline phones connect with an *RJ11* connector. The connector used with UTP cable is called *RJ45*. The RJ11 has room for two pairs (four wires), and the RJ45 has room for four pairs (eight wires). A crimper is used to attach an RJ connector to a cable. Higher-quality crimping tools

have interchangeable dies for both types of connectors. (Crimpers are discussed in Volume 1 Chapter 12, "Core Hardware and Network Troubleshooting.") Figure 6.26 shows two RJ11 and three RJ45 connectors. All have clips to keep them securely connected in the socket. The RJ45 connector on the right is from a shielded Cat 8 cable. The connector has visible gold shielding on three sides of it, whereas the standard RJ45 connectors next to it have no shielding.

FIGURE 6.26 RJ11 and RJ45 connectors

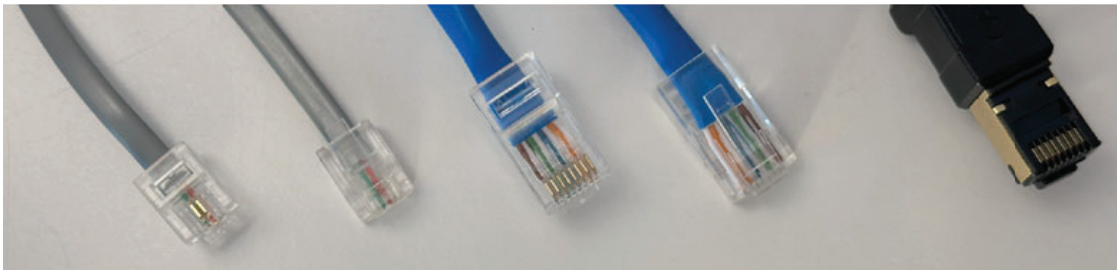

You will also find RJ45 splitters (often called Ethernet splitters) in the marketplace. The idea is similar to a coax splitter, but functionally they are very different. Coax signals are carried over one wire, while twisted pair uses either two pairs of wires (for 100 Mbps or slower connections) or all four pairs of wires (for Gigabit Ethernet and faster). An Ethernet splitter will take the incoming signal on two pairs and then split it, so on the output end it produces two sets of signals using two pairs each. Because of this, Ethernet splitters are limited to 100 Mbps connections.

It is not recommended that you use Ethernet splitters on a network. If you need to connect multiple computers together using UTP, use a hub or a switch. We talk about both of these devices later in this chapter.

Punchdown Blocks

Some twisted pair installations don't use standard RJ45 connectors. Instead, the cable is run to a central panel called a *punchdown block*, often located in a server room or connectivity closet. In a punchdown block, the metal wires are connected directly to the block to make the connection. Instead of a crimper, a punchdown tool is used. Figure 6.27 shows a closeup of wires connected to an older-style 66 block, frequently used in analog telephone communications. Networks that use blocks today are more likely to use a 110 block, which has a higher density of connectors and is designed to reduce crosstalk between cables.

Punchdown blocks are part of Exam Objective 3.2. So are T568A/T568B standards and direct burial, covered in the next two sections.

FIGURE 6.27 Cables in a punchdown block

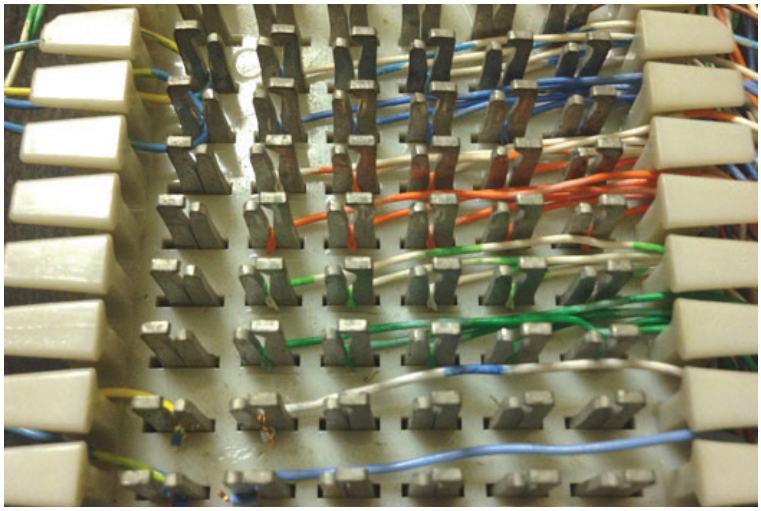

By Z22 - Own work, CC BY-SA 4.0, https://commons.wikimedia.org/w/index.php?curid=34324171

Wiring Standards

Twisted pair cables are unique in today's network environment in that they use multiple physical wires. Those eight wires need to be in the right places in the RJ45 connector or it's very likely that the cable will not work properly. To ensure consistency in the industry, two standards have been developed: *T568A* and *T568B*.

Older implementations using UTP used only two pairs of wires, and those two pairs were matched to pins 1, 2, 3, and 6 in the connector. Newer applications such as Voice over IP (VoIP) and Gigabit Ethernet use all four pairs of wires, so you need to make sure that they're all where they're supposed to be.

If you're creating a regular network *patch cable* to connect a computer to a hub or switch, both sides need to have the same pinout. For that, follow either the T568A standard shown in Figure 6.28 or the T568B standard shown in Figure 6.29. Although there are no differences in terms of how the standards perform, some companies prefer one to the other. (T568B seems to be more common.)

If you are going to create a cable to connect a computer to another computer directly, or you're going to make a connection from hub to hub, switch to switch, hub to switch, or a computer directly to a router, then you need what's called a *crossover cable*. In a crossover cable, pin 1 to pin 3 and pin 2 to pin 6 are crossed on *one side of the cable only*. This is to get the "send" pins matched up with the "receive" pins on the other side, and vice versa. For easier visualization, look at Figure 6.29.

The key thing to remember is that a patch (straight-through) cable is the same on both ends. A crossover cable is different on each end. You should know the order of the colors for both standards.

FIGURE 6.28 T568A standard

Pin	Pair	Wire	Color
1	3	1	white/green
2	3	2	green
3	2	1	white/orange
4	1	2	blue
5	1	1	white/blue
6	2	2	orange
7	4	1	white/brown
8	4	2	brown

FIGURE 6.29 T568B standard

Pin	Pair	Wire	Color
1	2	1	white/orange
2	2	2	orange
3	3	1	white/green
4	1	2	blue
5	1	1	white/blue
6	3	2	green
7	4	1	white/brown
8	4	2	brown

Direct Burial

Occasionally you will run into situations where network cable needs to be run outside or buried underground. For these types of installations, use *direct burial* cable. Direct burial cable is STP with an extra waterproof sheathing.

Whenever you run cables in an area where they can be stepped on (and it's not recommended you do), be aware that no amount of shielding will totally protect the cable from damage. It's best to use a cable guard of some sort to provide protection. We can't count the number of times we've seen people use duct tape or something similar to keep a cable in a high-traffic area from moving around—don't do it. The tape will work to keep it in place but does nothing to protect the cable.

An alternative may be to bury the cable underground. The recommended distance is 6″ to 8″ below the ground, and away from any lines that carry electrical current. Also, it's recommended that you use a conduit, such as PVC pipe, to further protect the cable.

Fiber-Optic Cable

Fiber-optic cabling offers a huge upgrade over copper cabling, providing transmission speeds from 100 Mbps to 10 Gbps over a maximum distance of several kilometers. It consists of a thin, flexible glass or plastic fiber surrounded by a rubberized outer coating. Because it uses pulses of light instead of electric voltages to transmit data, it is immune to electrical interference and to wiretapping.

Similarly to how we showed you the insides of a copper cable earlier, Figure 6.30 shows a cut fiber-optic cable. Admittedly the photo is grainy because we had to zoom in so much to show details (we're trainers, not professional photographers!), but what Figure 6.30 is attempting to show is the outer jacket, strength members (the Kevlar strings), and the coating and core of the cable. It might be hard to see, but there's a tiny dot in the center of the long tube of coating. That's the fiber cable itself. It's very, very thin.

FIGURE 6.30 Fiber-optic cable

Recognizing that Figure 6.30 might be a little hard to see, we had a professional draw what a fiber-optic cable looks like inside in Figure 6.31. And for good measure, a fiber-optic cable is shown next to an Ethernet Cat 8 cable for comparison in Figure 6.32. Fiber is thinner and more flexible than copper, but if its bent too much the cable can be ruined.

While it's gaining ground rapidly, fiber-optic cable is still not as popular as UTP for local area networks, however, because of its high cost of installation. Fiber-optic cabling is great for networks that need extremely fast transmission rates or transmissions over long distances or in networks that have had problems with electrical interference in the past. Fiber is also common as backbones to telecommunication systems, and in many places fiber-optic cables can be used to deliver high-speed Internet connections to businesses and homes. We'll talk more about this in Volume 1 Chapter 9.

FIGURE 6.31 Drawing of a fiber-optic cable

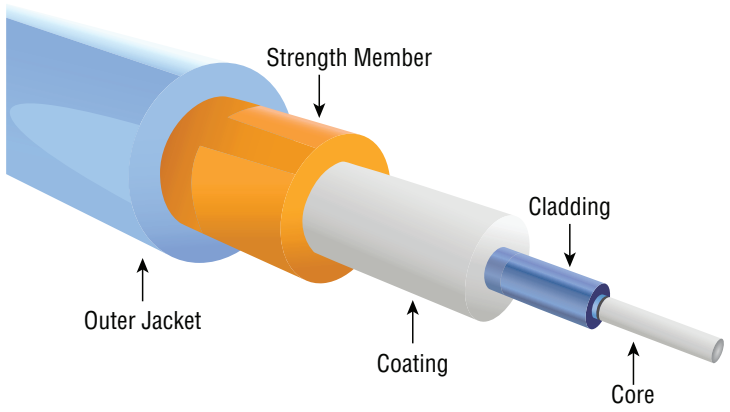

FIGURE 6.32 Ethernet Cat 8 (top) and fiber-optic cables (bottom)

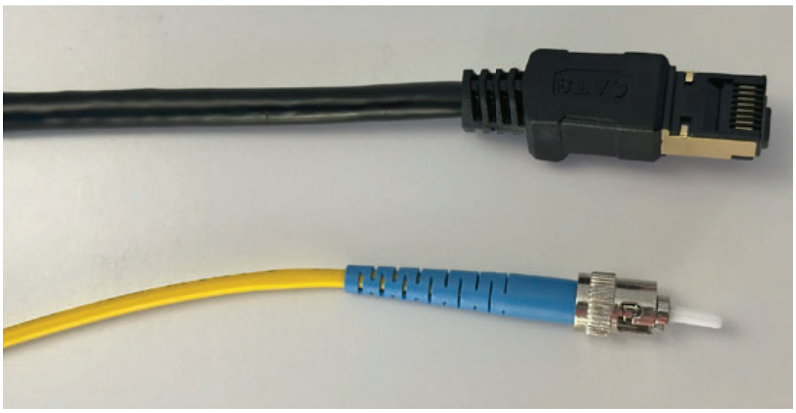

Fiber-Optic Cable Specifications

Fiber-optic cable comes in two varieties: single-mode or multimode. The term *mode* refers to the bundles of light that enter the fiber-optic cable. *Single-mode fiber (SMF)* cable uses only a single mode (or path) of light to propagate through the fiber cable, whereas *multimode fiber (MMF)* has a larger core that allows multiple modes of light to propagate simultaneously. In multimode fiber-optic cable, the light bounces off the cable walls as it travels through the cable, which causes the signal to weaken more quickly.

Single-mode and multimode optical cables are listed in Exam Objective 3.2. Be sure to understand the differences for the exam.

Multimode fiber is most often used as horizontal cable—that is, the cable between the backend infrastructure and the end-user systems. It permits multiple modes of light to propagate through the cable, which shortens cable distances but delivers more available

bandwidth. Devices that use MMF cable typically use light-emitting diodes (LEDs) to generate the light that travels through the cable; however, lasers with multimode fiber-optic cable are now being used in higher-bandwidth network devices, such as Gigabit Ethernet. MMF can transmit up to 400 Gbps for up to about 150 meters (492 feet) or lower data speeds such as 1 Gbps for up to 2 kilometers (about 1.2 miles), depending on the quality of cable and standard used.

Single-mode fiber cable is commonly used as backbone cabling, connecting telecommunications and equipment (or server) rooms together. It is also usually the cable type used in phone systems. Light travels through single-mode fiber-optic cable using only a single mode, meaning that it travels straight down the fiber and does not bounce off the cable walls. Because only a single mode of light travels through the cable, single-mode fiber-optic cable supports lower bandwidth at longer distances than does multimode fiber-optic cable. Devices that use single-mode fiber-optic cable typically use lasers to generate the light that travels through the cable. SMF can transmit up to 100 Gbps for up to 40 kilometers (25 miles) or slower speeds such as 1 Gbps up to 140 kilometers (87 miles), depending on the quality of cable and standard used.

Manufacturers grade fiber-optic cables based on quality. Single-mode fiber will have a grade such as *optical singlemode 1 (OS1)* or OS2. The higher the number, the more data it can carry over longer distances. Multimode fiber will have an *optical multimode (OM)* designation. Current OM ratings go from OM1 to OM5. Like with OS ratings, higher numbers in OM ratings indicate more capable cables.

We have talked about several different types of cables, and it's possible that you will be asked to know maximum distances and transmission speeds on the A+ exam. Table 6.4 summarizes some common cable types, the specifications with which they are used, and their characteristics.

TABLE 6.4 Common cable types and characteristics

Cable type	Ethernet specification	Maximum speed	Maximum distance	Notes
RG-6 coax	[a]	[a]	304 meters	Digital cable/satellite television.
RG-59 coax	[a]	[a]	228 meters	Analog cable TV.
Cat 5 UTP or STP	100BaseT	100 Mbps	100 meters	100 Mbps and less use two pairs of wires.
Cat 5e UTP	1000BaseT	1 Gbps	100 meters	1 Gbps and higher use four pairs of wires.
Cat 6 UTP	10GBaseT	10 Gbps	55 meters	Can support 1 Gbps up to 100 meters.

TABLE 6.4 JUNOS software Preference Values *(continued)*

Cable type	Ethernet specification	Maximum speed	Maximum distance	Notes
Cat 6a UTP	10GBaseT	10 Gbps	100 meters	
Cat 7 UTP	10GBaseT	10 Gbps	100 meters	Every wire pair is individually shielded.
CAT 8 UTP	25GBaseT or 40GBaseT	40 Gbps	100 meters at 10 Gbps	25 Gbps or 40 Gbps at 30 meters.
MMF fiber	1000BaseLX or 1000BaseSX	1 Gbps	550 meters	For fiber, maximum length depends on fiber size and quality.
MMF fiber	10GBaseSR or 10GBaseSW	10 Gbps	300 meters	
MMF fiber	100GBaseSR4	100 Gbps	100 meters	100 meters using OM4.
SMF fiber	10GBaseER or 10GBaseEW	10 Gbps	40 kilometers	
SMF fiber	100GBaseER4	100 Gbps	40 kilometers	

[a] RG-6 and RG-59 coax cables can be used with many different specifications, and the maximum speed depends on cable quality and specification.

So Many Standards!

Looking at Table 6.4 might give you some anxiety—how does one manage to memorize so many standards? While there are a lot of standards out there, you can usually discern what they mean if you follow the patterns.

Earlier in the chapter we mentioned that the number before the standard is the maximum speed, and that the word Base means baseband communication. This holds true even for 100 Gbps or 400 Gbps standards. After the number, if there is the letter T, the standard is for twisted pair. In general, think of twisted pair as having a maximum distance of 100 meters, although there are a few notable exceptions where it's less than that. It's never more than 100 meters, though.

If after Base there are other letters besides T, it must be fiber-optic. This is where things get a bit more complicated because there are a lot of fiber standards out there. Table 6.5 lays

them out. In general, though, think of MMF as higher bandwidth and shorter distance, up to around 500 meters, and SMF as the longest-distance option.

The good news is that it's highly unlikely that you will get asked about these designations on the exam, but they are helpful to know if you are dealing with fiber-optic cabling and components. If you want to memorize the order in Table 6.5, try the mnemonic, "Very short distance for long extended zeros." The first letter also means the same thing across fiber standards. For example, there is 10GBaseSR and 10GBaseSW. Both are short reach because they have an S. The SR specification calls for a standard Ethernet fiber connection, whereas SW specifies a network transmission called synchronous optical network (SONET), which is a competing transmission method to Ethernet. (You don't need to memorize that; it was just an example.)

TABLE 6.5 Fiber-optic letter designations

Abbreviation	Standard	Max distance	Typical cable type
VR	Very short reach	100 m	MMF
SR	Short reach	300 m	MMF
DR	Distance reach	500 m	MMF/SMF
FR	Far reach	2 km	SMF
LR	Long reach	10 km	SMF
ER	Extended reach	40 km	SMF
ZR	Zero-dispersion reach	80 km to 140 km	SMF

Finally, some of the newer standards have a number at the end, such as 400GBase-SR4 and 400GBase-SR8. The number at the end designates how many channels are used to achieve the maximum transmission speed. 400GBase-SR4 uses 4 parallel channels at 100 Gbps each to get to 400 Gbps, whereas 400Gbase-SR8 uses 8 channels at 50 Gbps each.

The key point to all of this is, please don't spend hours trying to memorize all of these standards. The goal is to equip you so that if you see a designation you're not familiar with, such as "100GBase-EX4," you can piece it out . . . okay, it's 100 Gbps, really long-range (extended reach), probably SMF, and uses 4 channels. And, know that if the equipment on one side of the connection is using that standard, the cable and the equipment on the other side of the connection need to support it as well.

Fiber-Optic Connector Types

There are literally dozens of fiber-optic connectors out there because it seemed that every producer wanted its proprietary design to become "the standard." Three of the most commonly used ones are ST, SC, and LC.

Straight tip (ST), subscriber connector (SC), and Lucent connector (LC) are all subobjectives of Exam Objective 3.2. Be able to identify these connector types if presented with one on the exam.

The *straight tip (ST)* fiber-optic connector, developed by AT&T, is probably the most widely used fiber-optic connector. It uses a twist-and-lock attachment mechanism that makes connections and disconnections fairly easy. The ease of use of the ST is one of the attributes that make this connector so popular.

The *subscriber connector (SC)* is also sometimes known as a *square connector*. SCs are latched connectors, making it virtually impossible for you to pull out the connector without releasing its latch, usually by pressing a button or release. SCs work with either single-mode or multimode optical fibers. They aren't as popular as ST connectors for LAN connections.

The last type of connector with which you need to be familiar is the *Lucent connector (LC)*, sometimes also called a *local connector*, which was developed by Lucent Technologies. It is a mini form factor (MFF) connector, especially popular for use with Fibre Channel adapters, fast storage area networks, and Gigabit Ethernet adapters. All three connectors are shown in Figure 6.33. Often you will see two connectors together in a pair (like the SC and LC connectors in Figure 6.33) so that the communication between two devices can be full-duplex.

FIGURE 6.33 ST (left), SC (center), and LC (right) fiber connectors

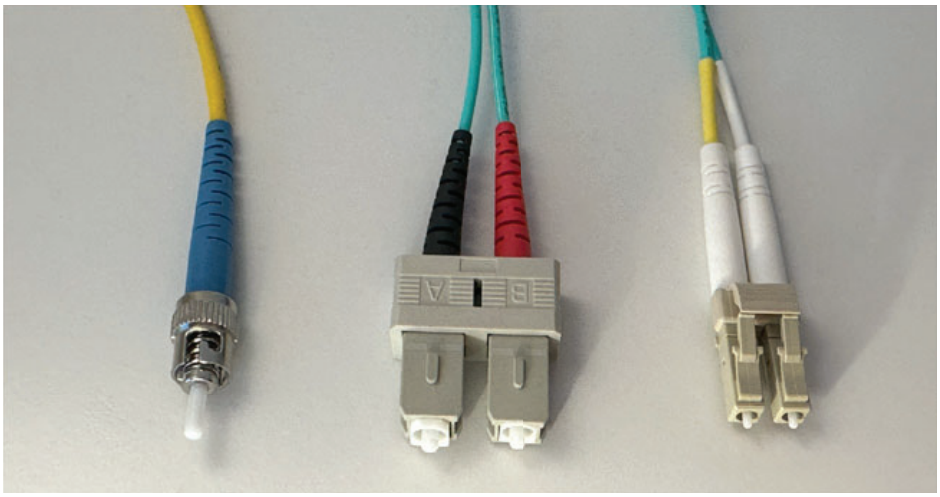

The prices of network cables differ dramatically between copper and fiber cables. Exercise 6.1 asks you to investigate the difference for yourself.

EXERCISE 6.1

Pricing Network Cables

1. Visit a major electronics retailer website of your choice or use an Internet search.
2. Search for a Cat 6 patch cable.
3. Price the difference between a 7-foot, 25-foot, and 50-foot cable.
4. Search for the same lengths of Cat 7 patch cables.

   Note the price difference. (As of this writing, Cat 7 cables cost a few dollars more than Cat 6 cables.)
5. Search for fiber-optic cables.

Notice, first, that most of them are much shorter in length than commercially available UTP cables. What is the price difference? Do you notice price differences between fiber-optic cables with different types of connectors?

Networking Components

Network cabling can link one computer to another, but most networks are far grander in scale than two simple machines. There are a variety of networking devices that provide connectivity to the network, make the network bigger, and offer auxiliary services to end users.

In the following sections, we're going to classify additional networking components into two broad categories: connectivity devices and auxiliary devices.

Connectivity Devices

We all know that if you want to be part of a computer network, you need to attach to that network somehow. Using network cables is one way to accomplish this, but not everyone is in a position to just plug a cable in and go. In addition, if you want to grow your network beyond a few simple connections, you need to use a special class of networking devices known as *connectivity devices*. These devices allow communications to break the boundaries of local networks and really provide the backbone for nearly all computer networks, regardless of size.

There are several categories of connectivity devices. These connectivity devices make it possible for users to connect to networks and to lengthen networks to almost unlimited distances. We will now discuss the most important and frequently used connectivity devices.

Exam Objective 2.5 lists several networking hardware devices that you should know about. They include cable modem, digital subscriber line (DSL), optical network terminal (ONT), access points, patch panels, switches (managed and unmanaged), and routers.

Modems

Let's start off by reminding you that modems are no longer part of the A+ exam objectives. However, RJ11 connectors are. And the only device we know of that uses RJ11 connectors is a modem.

If you want to connect to a network or the Internet using plain old phone lines and a dial-up connection, a *modem* is the device you'll need. Modems got their name because they modulate and demodulate (mo-dem) digital signals that computers use into analog signals that can be passed over telephone lines. In the early to mid-1990s, modems were practically the only device available to get onto the Internet. Many companies also used them to allow users who were not in the office to dial into the local network.

While modems did provide flexibility, you needed to be near a phone line, and speed was an issue. The fastest modems transferred data at 56 Kbps. At the time that felt lightning quick, but fortunately our species has moved well beyond that technology. It's horrifically slow by today's standards and therefore rarely used.

Cable/Digital Subscriber Line (DSL) Modems

The traditional modem is obsolete—accessing the Internet through the use of a *cable modem* or *digital subscriber line (DSL) modem* replaced them. The primary difference between the two is the infrastructure they connect to. Cable modems use television cable lines, and DSL modems use telephone lines.

Both cable and DSL modems are digital and therefore aren't technically modems because they don't modulate and demodulate analog signals. We'll cover cable Internet and DSL technologies and setting up a network using one of these devices in more detail in Volume 1 Chapter 9. Figure 6.34 shows a DSL modem. Cable modems are nearly identical, except they have an F-type connector for the coaxial cable connection.

Optical Network Terminal (ONT)

Fiber-optic connections to businesses and homes are becoming more and more common, as communications' providers race to install fiber all over the country. If there is fiber in your work or home neighborhood, you need a different type of device to connect to the ISP for Internet access. Such a device is called an *optical network terminal (ONT)*.One is shown in Figure 6.35.

Much like cable and DSL modems, an ONT isn't truly a modem either, as it doesn't deal with analog-to-digital modulation. It is closer to a modem in a sense, though, because it takes optical signals and changes them into electrical ones for your internal home or business network. ONTs are typically located out of sight in a wiring closet or at the junction box on the outside of the building, where the optical cabling comes to an end.

Access Points

Technically speaking, an *access point* is any point that allows a user on to a network. On a wired network, this means a hub or a switch, both of which we will cover shortly. The term is commonly used in reference to a *wireless access point*, which lets users connect to your network via an 802.11 technology. We'll get deeper into wireless access points and how to

configure them in Volume 1 Chapter 9. The back of a Linksys wireless router (which serves as a wireless access point) is shown in Figure 6.36.

FIGURE 6.34 DSL modem

Repeaters and Extenders

A *repeater*, or *extender*, is a small, powered device that receives a signal, amplifies it, and sends it on its way. The whole purpose of a repeater is to extend the functional distance of a cable run. For example, you know that UTP is limited to 100 meters, but what if you need to make a cable run that is 160 meters long? (One answer could be to use fiber, but pretend that's not an option.) You could run two lengths of cable with a repeater in the center, and it would work. Repeaters and extenders work at the Physical layer (Layer 1) of the OSI model. They don't examine the data or make any changes to it—they just take what they receive and send it along its merry way.

Hubs

A *hub* is a device used to link several computers together. Hubs are very simple devices that possess no real intelligence. They simply repeat any signal that comes in on one port and copy it to the other ports (a process that is also called *broadcasting*). You'll sometimes hear them referred to as multiport repeaters. They work at Layer 1 of the OSI model, just as repeaters do.

FIGURE 6.35 Optical network terminal (ONT)

FIGURE 6.36 Wireless router

There are two types of hubs: active and passive. *Passive hubs* connect all ports together electrically but do not have their own power source. Think of them as a multiport repeater. *Active hubs* use electronics to amplify and clean up the signal before it is broadcast to the other ports.

Active hubs can therefore be used to extend the length of a network, whereas passive hubs cannot. Hubs are largely obsolete today, replaced by switches. You will hear people refer to switches as hubs, though, because they are a central point of connection. The two are fundamentally different though, as you will learn in the upcoming "Switches" section.

Patch Panels

A *patch panel* is essentially a large hub that is rack mounted. It houses multiple cable connections but possesses no network intelligence. Its sole purpose is to connect cables together. Short patch cables are used to plug into the front-panel connectors, and there are longer, more permanent cables on the back. Figure 6.37 shows several rack-mounted patch panels in a server room.

FIGURE 6.37 Rack-mounted patch panels

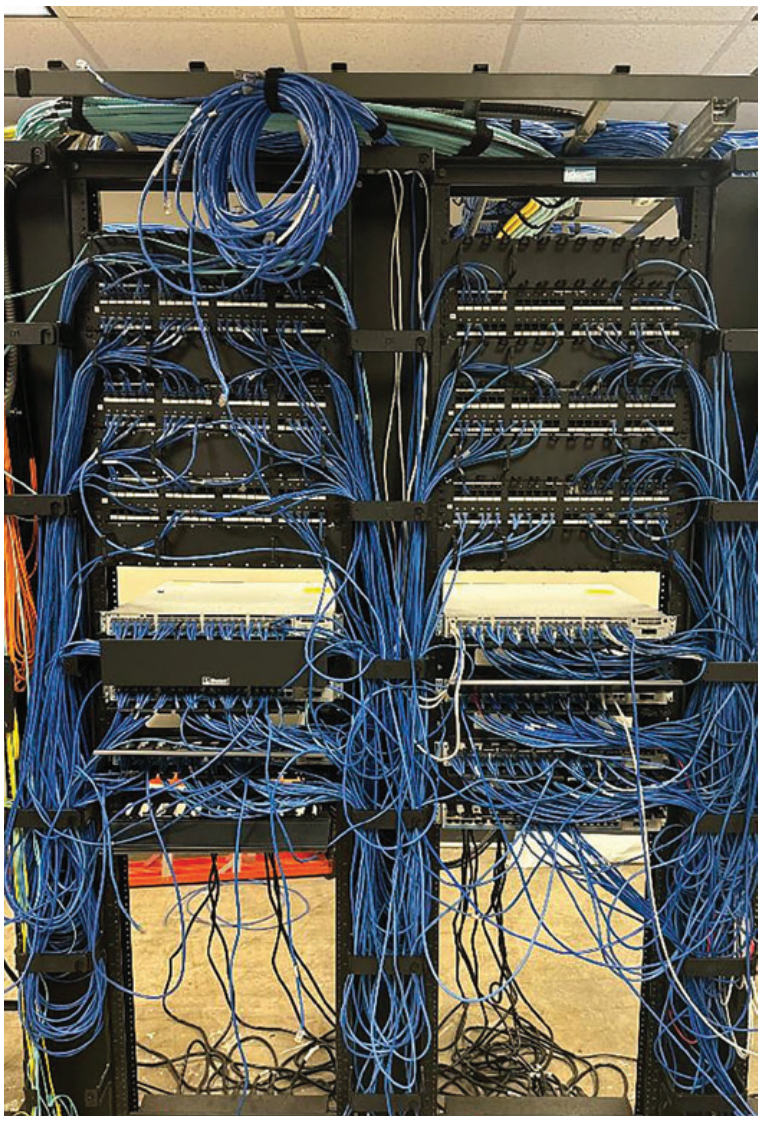

Switches

Switches work at Layer 2 (unless it's a multilayer switch, and then it works at Layer 3 as well—but ignore that for now) and they provide centralized connectivity, just like hubs. A *switch* often looks similar to a hub, so it's easy to confuse them. There are big performance differences, though. Hubs pass along all traffic, but switches examine the Layer 2 header of the incoming packet and forward it properly to the right port and only that port. This greatly reduces overhead and thus improves performance because there is essentially a virtual connection between sender and receiver. The only downside is that switches forward broadcasts because they are addressed to everyone.

Switches come in two varieties: unmanaged and managed. We've already explained the functionality of an *unmanaged switch*—it connects two or more computers, and passes along all traffic sent to a MAC address to its port. A *managed switch* adds the ability to configure ports, manage traffic, and monitor traffic for issues. For management, the switch will use a network protocol, such as Simple Network Management Protocol (SNMP). (We'll talk about SNMP in depth in Volume 1 Chapter 7.) Managed switches cost more but provide features such as quality of service (QoS), redundancy, port mirroring, and virtual LANs (VLANs). Here's a description of each:

QoS QoS allows administrators to prioritize certain network traffic by assigning a higher priority to it. Higher-priority traffic may come from a specific server or a specific application. This is used a lot with Voice over IP (VoIP)—telephone calls over a computer network—to ensure that the voice data gets through and the connection doesn't sound garbled.

Redundancy Redundancy in networking terms means having multiple paths to get data from point A to point B. Administrators can use multiple switches to provide redundant paths, which add a layer of fault tolerance to the network. Managed switches use the Spanning Tree Protocol (STP) to implement redundancy.

Port Mirroring This is a troubleshooting feature that is used in conjunction with a network monitor. A port can be configured to mirror another port. When traffic is sent to one, it's also sent to the mirror. A network monitor attached to the mirrored port can then analyze the traffic, without taking the network or device on the original port offline.

VLANs In a virtual LAN (VLAN), computers attached to the same physical switch can be segmented into multiple logical networks. This reduces network traffic on each VLAN, because the traffic is isolated from other virtual LANs. Computers on one virtual LAN can still communicate with those on another VLAN, if the switch is configured properly. VLANs can also be completely isolated from each other, adding an additional level of security.

Nearly every hub or switch that you will see has one or more status indicator lights on it. If there is a connection to a port of the switch, a light either above the connector or on an LED panel elsewhere on the device will light up. If traffic is crossing the port, the light may flash, or there may be a secondary light that will light up. Many devices can also detect a

problem in the connection. If a normal connection produces a green light, a bad connection might produce an amber light.

Small switches for home or small office network use might have only 6 or 8 ports. Larger rack-mounted switches are found in server rooms of bigger networks. Figure 6.38 shows two Cisco switches below two Cisco 4300 series routers.

FIGURE 6.38 Cisco routers and switches

While it's common practice to refer to switches as Layer 2 devices, some do work at Layer 2 and Layer 3. They are called multilayer switches. A Layer 2 switch concerns itself with MAC addresses only and doesn't pay attention to IP addresses. A Layer 3 switch will also work with IP addresses and can handle routing functionality between VLANs.

Routers

Routers are highly intelligent devices that connect multiple network types and determine the best path for sending data. They can route packets across multiple networks and use *routing tables* to store network addresses to determine the best destination. Routers operate at the Network layer (Layer 3) of the OSI model. Because of this, they make their decisions on what to do with traffic based on logical addresses, such as an IP address.

Routers have a few key functions:

- They connect multiple networks to each other, which none of the other devices we have discussed do.
- Routers do not forward broadcasts. (Switches and bridges break up collision domains, whereas routers break up broadcast domains.)
- Routers are normally used to connect one LAN to another. Typically, when a WAN is set up, at least two routers are used.

Two rack-mounted Cisco routers are shown in Figure 6.38. In the last decade or so, small wireless routers (like the one in Figure 6.36) have become common for small business and home

networks. They possess all the functionality of routers historically associated with networking, but they are relatively inexpensive. We'll talk more about these routers in Volume 1 Chapter 9. An example network using switches, routers, and a firewall is shown in Figure 6.39.

FIGURE 6.39 Network example

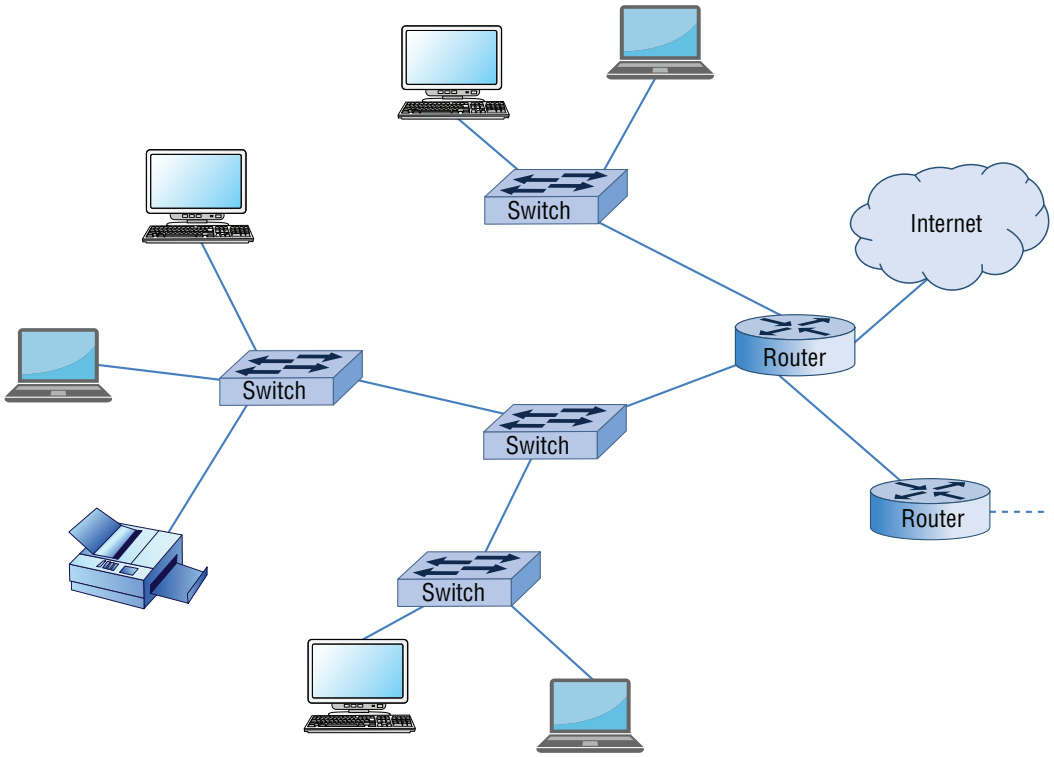

Auxiliary Devices

The devices we just talked about are specialized to provide connectivity. This next group of devices adds in features outside of connectivity that can help network users, specifically by protecting them from malicious attacks or providing power over Ethernet cables.

Firewalls and Power over Ethernet standards, injectors, and switches are part of Exam Objective 2.5.

Firewall

A *firewall* is a hardware or software solution that serves as your network's security guard. They're probably the most important devices on networks that are connected to the Internet. Firewalls can protect you in two ways: They protect your network resources from hackers

lurking in the dark corners of the Internet, and they can simultaneously prevent computers on your network from accessing undesirable content on the Internet. At a basic level, firewalls filter packets based on rules defined by the network administrator.

Firewalls can be stand-alone "black boxes," software installed on a server or router, or some combination of hardware and software. Most firewalls will have at least two network connections: one to the Internet, or *public side*, and one to the internal network, or *private side*. Some firewalls have a third network port for a second semi-internal network. This port is used to connect servers that can be considered both public and private, such as web and email servers. This intermediary network is known as a *screened subnet* (formerly called a *demilitarized zone [DMZ])*.

A screened subnet is a semi-public network segment located between a perimeter router and an internal router on your network. It is used for web servers, FTP servers, and email relay servers.

Firewalls can be network-based in that they protect a group of computers (or an entire network), or they can be host-based. A host-based firewall (such as Windows Defender Firewall) protects only the individual computer on which it's installed.

A firewall is configured to allow only packets that pass specific security restrictions to get through. By default, most firewalls are configured as *default deny*, which means that all traffic is blocked unless specifically authorized by the administrator. The basic method of configuring firewalls is to use an *access control list (ACL)*. The ACL is the set of rules that determines which traffic gets through the firewall and which traffic is blocked. ACLs are typically configured to block traffic by IP address, port number, domain name, or some combination of all three.

We'll cover firewalls in more depth in Volume 1 Chapter 9 when we show you how to set up a network.

Power over Ethernet

Imagine a situation where you need to have a network device such as a wireless access point or IP security camera in a relatively remote location where there are no power outlets. Installing new outlets might be impractical or expensive, so there is a helpful solution called *Power over Ethernet (PoE)*.

There are two ways to get power to a device through an Ethernet cable. If you only need to power a single device, then a *Power over Ethernet injector* (Figure 6.40) will do the trick. As you can see in Figure 6.40, it only has two connectors: one for the inbound data connection (from a router), and one to the endpoint device giving it power and a network connection.

FIGURE 6.40 Power over Ethernet injector

If there are multiple devices that need power, use a switch that's PoE enabled, such as the one shown in Figure 6.41. Simply plug a PoE-enabled endpoint device into ports 1–4 with an Ethernet cable, and the endpoint device will have power. Then using port 5, plug the PoE switch into the network router to enable connectivity. PoE switches cost more than injectors, but being able to power multiple devices at once is a bonus.

FIGURE 6.41 Power over Ethernet switch

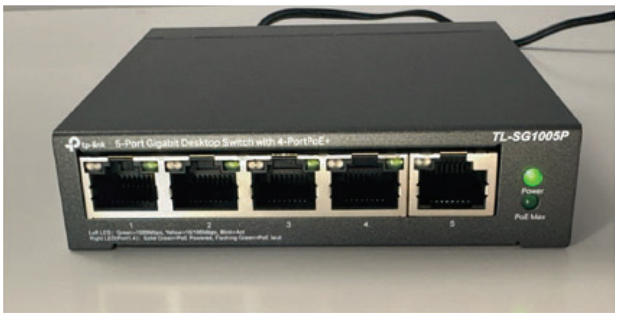

The first PoE standard was IEEE 802.3af, released in 2003, and it provided up to 15.4 W of DC power to connected devices. This was enough for wireless access points as well as basic surveillance cameras and VoIP phones, but not enough for videoconferencing equipment, alarm systems, laptops, or flat-screen monitors. Enhancements to the standard have been made over the years to support more power-hungry devices. Table 6.6 lists the standards you should be familiar with.

TABLE 6.6 PoE standards

Name	Year	IEEE standard	Max power	Supported devices
PoE	2003	802.3af	15.4 W	Wireless access points, static surveillance cameras, VoIP phones
PoE+	2009	802.3at	30 W	Alarm systems, PTZ (point/tilt/zoom) cameras, video IP phones
PoE++	2018	802.3bt (Type 3)	60 W	Multi-radio wireless access points, video conferencing equipment
PoE++	2018	802.3bt (Type 4)	100 W	Laptops, flat-screen monitors, LED lighting

Summary

In this chapter, we covered a broad variety of networking topics. This chapter contains everything that you need to get you ready for the networking questions on the A+ 220-1201 exam. At the same time, the A+ exam (and consequently this chapter) barely scratches the surface of the things that you can learn about networking. If making computers talk to each other effectively is an area of interest to you, we suggest that you consider studying for the CompTIA Network+ exam after you pass your A+ tests.

First, we started with networking fundamentals. Much of the discussion of fundamentals was about understanding the concepts behind networking so that you know how to set them up. Topics included LANs versus WANs; clients, servers, and resources; network operating systems; peer-to-peer and server-based resource models; network topologies, such as bus, star, and ring; and theoretical networking models and standards, such as the OSI model and IEEE standards.

Next, you learned about hardware devices used in networking. Each computer needs a network adapter (NIC) of some sort to connect to the network. On a wired network, cables are required, and there are several different types, including coaxial, STP, UTP, and fiber-optic. Each cable type has its own specific connector.

Finally, we discussed various types of network connectivity hardware and auxiliary devices and their use. Some users may need a cable modem, DSL modem, ONT, or access point to get onto the network. All wired computers will plug into a connectivity device, such as a switch, which in turn is connected to another connectivity device, which is often a router. Other devices on the network, such as firewalls and PoE injectors and switches provide additional services.

Exam Essentials

Know the difference between LANs, WANs, PANs, MANs, SANs, and WLANs. A LAN is a local area network, which typically means a network in one centralized location. A WAN is a wide area network, which means several LANs in remote locations connected to each other. A PAN is a small Bluetooth network. A network that spans an area such as a city or a campus is a MAN. A SAN is designed specifically for storage, and a WLAN is like a LAN but wireless.

Know how computers connect to a network. It might seem simple, but remember that all computers need a NIC to connect to the network. There's a lot of configuration that happens automatically, and you may need to reconfigure the NIC or update drivers if things don't work properly.

Understand what a MAC address is. The media access control (MAC) address is physically associated with a network card. MAC addresses must be unique for them to work on a network.

Know about the different types of copper cable. The three types of copper cable you should know about are coaxial, unshielded twisted pair (UTP), and shielded twisted pair (STP). UTP comes in various types, including Cat 5, Cat 5e, Cat 6, and Cat 6a (among others). For outdoor use, go with direct burial cable.

Understand the difference between a patch (straight-through) cable and a crossover cable. Patch cables are used to connect hosts to a switch or a hub. Crossover cables switch pins 1 and 3 and 2 and 6 on one end. They are used to connect hubs to hubs, switches to switches, hosts to hosts, and hosts to routers.

Memorize the T568A and T568B cable standards. As painful as it might sound, you should memorize the pin order for these two standards. The T568A order is white/green, green, white/orange, blue, white/blue, orange, white/brown, brown. T568B is white/orange, orange, white/green, blue, white/blue, green, white/brown, brown. If it helps, note that the blue and brown pairs do not change; only the green and orange pairs do.

Know what a plenum cable is used for. Plenum cables do not release toxic gas when burned and therefore are required in spaces that circulate air (plenums) within buildings.

Understand performance characteristics of fiber-optic cable. Fiber can support higher transmission rates and longer distances than copper cable can. It's also immune to electrical interference. Single-mode cables can run longer distances than multimode cables.

Know which types of connectors are used for the different types of network cables. Coaxial cable uses F-type or BNC connectors. Twisted pair uses RJ11 or RJ45 connectors or can be terminated at a punchdown block. Fiber connectors include straight tip (ST), subscriber connector (SC), and Lucent (or local) connector (LC).

Know which networking devices are used to connect to the Internet. Digital connections are made through cable modems and DSL modems, and optical connections through optical network terminals (ONTs).

Know what hubs, switches, access points, patch panels, and routers are. These are all network connectivity devices. Hubs and switches are used to connect several computers or groups of computers to each other. Switches can be managed or unmanaged. An access point is any port where a computer plugs into a network, but the term typically refers to wireless access points. Patch panels are rack-mounted devices with multiple (usually dozens of) wired access points. Routers are more complex devices that are often used to connect network segments or networks to each other.

Know what a firewall and Power over Ethernet (PoE) provides. A firewall is a security device that blocks or allows network traffic to pass through it. PoE provides for electricity over Ethernet cables.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. _____ is immune to electromagnetic or radio-frequency interference.
   - **A.** Twisted pair cabling
   - **B.** CSMA/CD
   - **C.** Broadband coaxial cabling
   - **D.** Fiber-optic cabling

2. Which of the following is a type of connector used with coaxial cable?
   - **A.** RJ11
   - **B.** RJ45
   - **C.** F-type
   - **D.** SC

3. You need to put a connector on the end of an Ethernet crossover cable. The existing end has a white and green colored wire at pin 1. What color do you need to make pin 1 on the new end?
   - **A.** White/green
   - **B.** Green
   - **C.** White/orange
   - **D.** Orange

4. _____ is the type of media access method used by NICs that listen to or sense the cable to check for traffic and send only when they hear that no one else is transmitting.
   - **A.** T568A/T568B
   - **B.** CSMA/CD
   - **C.** CSMA/CA
   - **D.** Demand priority

5. What is the lowest grade of UTP that is rated for 100 Mbps transmission?
   - **A.** Cat 5
   - **B.** Cat 5e
   - **C.** Cat 6
   - **D.** Cat 6a

6. An administrator needs to run a network cable between two buildings that are about 150 yards apart on a campus. Company regulations say the cable must be underground. What type of cable should the administrator use?
   - **A.** Fiber-optic
   - **B.** STP

C. Plenum

D. Direct burial

7. Of all the network cabling options, _____ offers the longest possible segment length.

A. Unshielded twisted pair

B. Coaxial

C. SMF

D. MMF

E. Shielded twisted pair

8. What devices transfer packets across multiple networks and use tables to store network addresses to determine the best destination?

A. Routers

B. Firewalls

C. Hubs

D. Switches

9. Which of the following networking hardware devices is capable of creating VLANs?

A. Firewall

B. Patch panel

C. Unmanaged switch

D. Managed switch

10. Which of the following wireless communication standards is often described in terms of a wireless personal area network?

A. Bluetooth

B. Infrared

C. Cellular

D. Ethernet

11. Your ISP has just finished running fiber-optic cable in your neighborhood, and you want to switch from cable Internet to a faster fiber-optic connection. Which of the following statements is true about how you will get an Internet connection?

A. You will continue to use your cable modem.

B. You need to switch to using a DSL modem.

C. You need to switch to using an ONT.

D. You need to switch to using PoE.

12. If you are going to run a network cable in the space above the drop ceiling in your office, which type of cable should you use?

A. Plenum

B. PVC

C. Coaxial

D. Fiber-optic

E. Direct burial

13. Which of the following connector types is an MFF connector?

A. BNC

B. ST

C. SC

D. LC

14. What Ethernet specification would you be running if you needed to make a connection of 10 Gbps over a distance of five kilometers?

A. 10GBaseER

B. 10GBaseT

C. 10GBaseSR

D. 10GBaseDR

15. Which of the following are advantages to using a SAN? (Choose two.)

A. More storage space on the network

B. More efficient data storage

C. Faster data access

D. Automated data backups

16. Which type of fiber-optic cable connector uses a twist to lock mechanism to secure the cable in place?

A. BNC

B. SC

C. ST

D. LC

17. What type of device will block unwanted traffic from your network using a set of rules called an ACL?

A. Router

B. Firewall

C. Switch

D. NAS

18. What type of coaxial cable is recommended for digital television cable signals?

A. RG-6

B. RG-8

C. RG-58

D. RG-59

19. Which type of device works with MAC addresses to determine what to do with network traffic?
    - **A.** Hub
    - **B.** Router
    - **C.** Patch panel
    - **D.** Switch

20. Which type of networking device uses bare wires for attachments rather than a connector?
    - **A.** Patch panel
    - **B.** Punchdown block
    - **C.** PoE injector
    - **D.** Access point

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answer compares to the authors', refer to Appendix B.

Look at the pictures of network cable connectors and label each one.

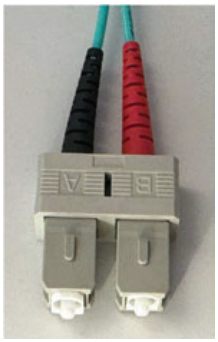

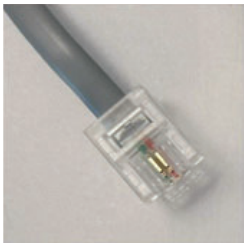

Chapter 7

Introduction to TCP/IP

THE FOLLOWING COMPTIA A+ 220-1201 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

✓ **2.0 Networking**

✓ **2.1 Compare and contrast Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) ports, protocols, and their purposes.**

- Ports and protocols
  - 20-21—File Transfer Protocol (FTP)
  - 22—Secure Shell (SSH)
  - 23—Telnet
  - 25—Simple Mail Transfer Protocol (SMTP)
  - 53—Domain Name System (DNS)
  - 67/68—Dynamic Host Configuration Protocol (DHCP)
  - 80—Hypertext Transfer Protocol (HTTP)
  - 110—Post Office Protocol 3 (POP3)
  - 143—Internet Mail Access Protocol (IMAP)
  - 137-139—Network Basic Input/Output System (NetBIOS)/NetBIOS over TCP/IP (NetBT)
  - 389—Lightweight Directory Access Protocol (LDAP)
  - 443—Hypertext Transfer Protocol Secure (HTTPS)
  - 445—Server Message Block (SMB)/Common Internet File System (CIFS)
  - 3389—Remote Desktop Protocol (RDP)
- TCP vs. UDP

✓ **2.4 Explain common network configuration concepts.**

- Virtual LAN [local area network] (VLAN)
- Virtual private network (VPN)

✓ **2.6 Given a scenario, configure basic wired/wireless small office/home office (SOHO) networks.**

- Internet Protocol (IP) addressing
  - IPv4
    - Private addresses
    - Public addresses
  - IPv6
  - Automatic Private IP Addressing (APIPA)
  - Static
  - Dynamic
  - Subnet mask
  - Gateway

Networking protocols are a lot like human languages in that they are the languages that computers speak when talking to each other. If computers don't speak the same language, they won't be able to communicate. To complicate matters, there are dozens of different languages that computers can use. Just like humans, computers can understand and use multiple languages. Imagine that you are on the street and someone comes up to you and speaks in Spanish. If you know Spanish, you will likely reply in kind. It doesn't matter if both of you know English as well because you've already established that you can communicate. On the other hand, it's going to be a pretty futile conversation if you don't know Spanish. This same concept applies to computers that are trying to communicate. They must have a network protocol in common in order for the conversation to be successful.

Throughout the years, hundreds of network protocols have been developed. As the use of networking exploded, various companies developed their own networking hardware, software, and proprietary protocols. Some were incorporated as an integral part of the network operating system, such as Banyan VINES. One-time networking giant Novell had IPX/SPX. Microsoft developed NetBEUI. Apple created AppleTalk. Others included DECnet, SNA, and XNS. While a few achieved long-term success, most have faded into oblivion. The one protocol suite that has survived is TCP/IP. Although it has some structural advantages, such as its modularity, it didn't necessarily succeed because it was inherently superior to other protocols. It succeeded because it is the protocol of the Internet.

This chapter focuses on the TCP/IP protocol suite. While it is the protocol suite used on the Internet, it's also the protocol suite used by nearly every home and business network today. We'll start by taking a quick look at the history of TCP/IP and the model on which it's based. Then we'll dive deeper into TCP/IP structure and the individual protocols it includes. From there, we'll spend some time on IP addressing, including IPv4 and IPv6. Entire books have been written on TCP/IP—so there's no way we could cover it entirely in one chapter. Nor do you need to know every last detail right now. Instead, we'll give you the foundation that you need to understand it well, work effectively with it in the field, and pass the A+ exam.

Understanding TCP/IP

As we mentioned in the introduction, computers use a protocol as a common language for communication. A *protocol* is a set of rules that govern communications, much like a language in human terms. Of the myriad protocols out there, the key ones to understand are the

protocols in the TCP/IP suite, which is a collection of different protocols that work together to deliver connectivity. Consequently, they're the only ones listed on the A+ exam objectives. In this section, we'll start with a look at its overall structure and then move into key protocols within the suite.

TCP/IP Structure

The *Transmission Control Protocol/Internet Protocol (TCP/IP) suite* is the most popular network protocol in use today, thanks mostly to the rise of the Internet. While the protocol suite is named after two of its hardest-working protocols, *Transmission Control Protocol (TCP)* and *Internet Protocol (IP)*, TCP/IP actually contains dozens of protocols working together to help computers communicate with one another. While it became popular due to the rise of the Internet, its current integration into cloud services and the Internet of Things (IoT) essentially guarantees it won't be going away any time soon.

TCP/IP is robust and flexible. For example, if you want to ensure that the packets are delivered from one computer to another, TCP/IP can do that. If speed is more important than guaranteed delivery, then TCP/IP can provide that too. The protocol can work on disparate operating systems, such as UNIX, Linux, macOS, Windows, iOS, and Android. It can also support a variety of programs, applications, and required network functions. Much of its flexibility comes from its modular nature.

You're familiar with the seven-layer OSI model that we discussed in Volume 1 Chapter 6, "Networking Fundamentals." Every protocol that's created needs to accomplish the tasks (or at least the key tasks) outlined in that model. The structure of TCP/IP is based on a similar model created by the U.S. Department of Defense—that is, the *Department of Defense (DoD) model*. The DoD model (sometimes referred to as the TCP/IP model) has four layers that map to the seven OSI layers, as shown in Figure 7.1.

FIGURE 7.1 The DoD and OSI models

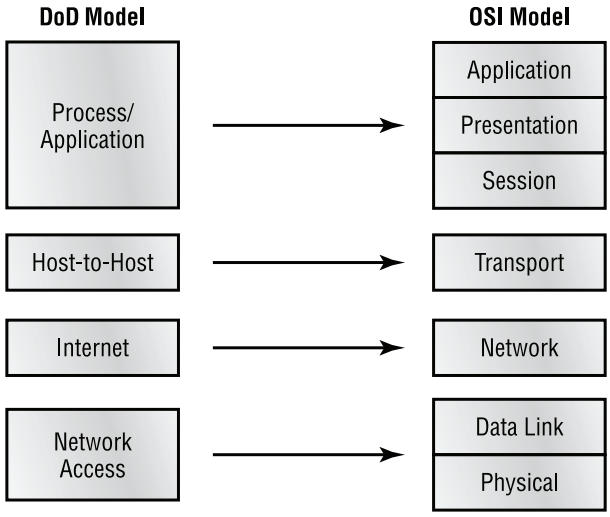

The overall functionality between these two models is virtually identical; the layers just have different names. For example, the Process/Application layer of the DoD model is designed to combine the functionality of the top three layers of the OSI model. Therefore, any protocol designed against the Process/Application layer would need to be able to perform all the functions associated with the Application, Presentation, and Session layers in the OSI model.

TCP/IP's modular nature and common protocols are shown in Figure 7.2.

FIGURE 7.2 TCP/IP protocol suite

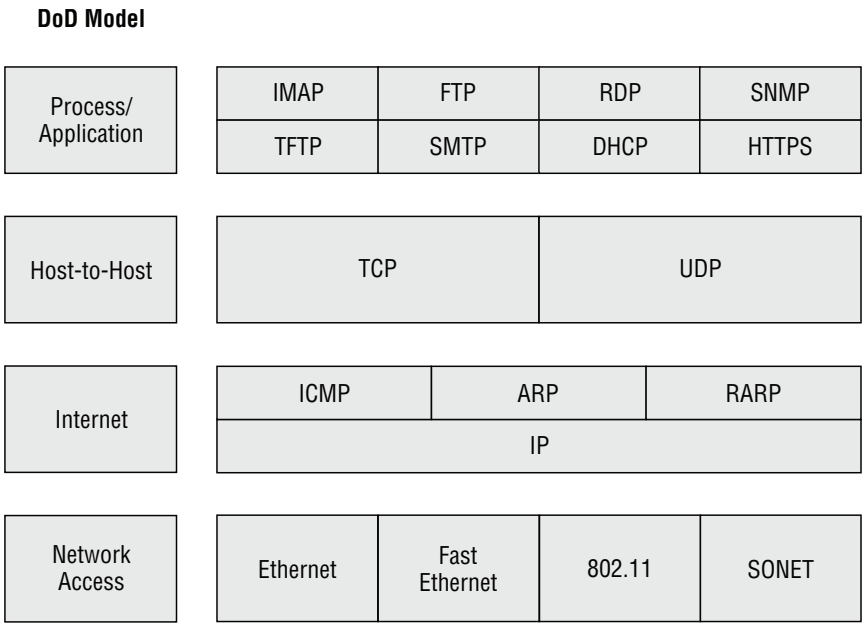

Working from the bottom up, you'll notice that the Network Access layer doesn't have any protocols, as such. This layer describes the type of network access method that you are using, such as Ethernet, Wi-Fi, or others.

The most important protocol at the Internet layer is IP. This is the backbone of TCP/IP. Other protocols at this layer work in conjunction with IP, such as *Internet Control Message Protocol (ICMP)* and *Address Resolution Protocol (ARP)*.

At the Host-to-Host layer, there are only two protocols: TCP and *User Datagram Protocol (UDP)*. Most applications will use one or the other to transmit data, although some can use both but will do so for different tasks.

The majority of TCP/IP protocols are located at the Process/Application layer. These include some protocols with which you may already be familiar, such as *Hypertext Transfer Protocol (HTTP)*, *Hypertext Transfer Protocol Secure (HTTPS)*, *File Transfer Protocol (FTP)*, *Simple Mail Transfer Protocol (SMTP)*, *Internet Message Access Protocol (IMAP)*, and others. Let's take a look at each of the layers in more detail.

Internet Layer Protocols

At the Internet layer, there's one key protocol and a few helpful support protocols. The main workhorse of TCP/IP is the *Internet Protocol (IP)*, and it can be found at this layer. IP is responsible for managing logical network addresses and ultimately getting data from point A to point B, even if there are dozens of points in between. We cover IP addressing in depth in the "Understanding IP Addressing" section later in this chapter.

There are three support protocols you should be aware of at this layer as well:

- *Internet Control Message Protocol (ICMP)* is responsible for delivering error messages. If you're familiar with the `ping` utility, you'll know that it utilizes ICMP to send and receive packets.
- *Address Resolution Protocol (ARP)* resolves logical IP addresses to physical MAC addresses built into network cards. This function is critical because in order to communicate, the sender ultimately needs to know the MAC address of the receiver.
- *Reverse ARP (RARP)* resolves MAC addresses to IP addresses.

Host-to-Host Layer Protocols

Next up is the Host-to-Host layer, and it has the fewest protocols. At this layer there are two alternatives within the TCP/IP suite: TCP and UDP. The major difference between the two is that TCP guarantees packet delivery through the use of a virtual circuit and data acknowledgments and UDP does not. Because of this, TCP is often referred to as *connection-oriented*, whereas UDP is *connectionless*. Because UDP is connectionless, it does tend to be somewhat faster, but we're talking about milliseconds here.

A+ Exam Objective 2.1, "Compare and contrast Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) ports, protocols, and their purposes," is central to this chapter. Specifically, be sure to understand the differences between TCP and UDP and why each has its place on a network. You will also need to know why port numbers are used, and memorize port numbers for common protocols.

Another key concept to understand about TCP and UDP is the use of *port numbers*. Imagine a web server that is managing connections from incoming users who are viewing web content and others who are downloading files. TCP and UDP use port numbers to keep track of these conversations and make sure that the data gets to the right application and right end user. Conversely, when a client makes a request of a server, it needs to do so on a specific port to make sure that the right application on the server hears the request. For example, web servers are listening for HTTPS requests on port 443, so web browsers need to make their requests on that port.

A good analogy for understanding port numbers is cable or satellite television. In this analogy, the IP address is your house. The cable company needs to know where to send the data. But once the data is in your house, which channel are you going to receive it on? If you want sports, that might be on one channel, but weather is on a different channel, and

the cooking show is on yet another. You know that if you want a cooking show, you need to turn to channel 923 (or whatever). Similarly, the client computer on a network knows that if it needs to ask a question in HTTPS, it needs to do it on port 443.

There are 65,536 ports, numbered from 0 to 65535. Ports 0 through 1023 are called the *well-known ports* and are assigned to commonly used services, and 1024 through 49151 are called the *registered ports*. All the ports from 49152 to 65535 are free to be used by application vendors. Fortunately, you don't need to memorize them all.

TCP/IP applications combine the host's IP address with the port number in order to communicate. This combination is known as a *socket*. You might see it written as something like `192.168.2.115:443`. Be aware that the term "socket" has the potential to appear on the A+ exam.

Table 7.1 shows the ports used by some of the more common protocols. You should know each of these for the A+ exam.

TABLE 7.1 Common port numbers

Service	Protocol	Port(s)
FTP	TCP	20, 21
SSH	TCP	22
Telnet	TCP	23
SMTP	TCP	25
DNS	TCP/UDP	53
DHCP	UDP	67, 68
HTTP	TCP	80
POP3	TCP	110
NetBIOS/NetBT	TCP	137-139
IMAP4	TCP	143
LDAP	TCP	389
HTTPS	TCP	443
SMB/CIFS	TCP	445
RDP	TCP	3389

A complete list of registered port numbers can be found at `https://iana.org` and several other sites, such as Wikipedia.

Process/Application Layer Protocols

As we mentioned earlier in the chapter, most of the protocols within the TCP/IP suite are at the Process/Application layer. This is the layer of differentiation and flexibility. For example, if you want to browse the Internet, the HTTP and HTTPS protocols are designed for that. FTP is optimized for file downloads, and Simple Mail Transfer Protocol (SMTP) is used for sending email.

Before we get into the protocols themselves, let's take a quick look into a few key points on the TCP/IP suite's flexibility. There are literally dozens of protocols at the Process/Application layer, and they have been created over time as networking needs arose. Take HTTP, for example. The first official version was developed in 1991, nearly 20 years after TCP/IP was first implemented. Before this protocol was created, there weren't any effective client-server request-response protocols at this layer. HTTP let the client (web browser) ask the web server for a page, and the web server would return it. Going one step further, there was a need for secure transactions over HTTP—hence, the creation of HTTPS in 1994. Now, almost all websites use HTTPS, and HTTP is fading away.

As new applications are developed or new networking needs are discovered, developers can build an application or protocol that fits into this layer to provide the needed functionality. They just need to make sure that the protocol delivers what it needs to and can communicate with the layers below it. The following sections describe some of the more common Process/Application protocols and their ports—and the ones listed in the A+ exam objectives.

It might be apparent based on our comments so far, but you should know each of the upcoming protocols, what they do, and what ports they use before attempting the A+ exam. They are all outlined in Exam Objective 2.1.

Port 20/21—File Transfer Protocol (FTP)

The *File Transfer Protocol (FTP)* is optimized to do what it says it does—transfer files. This includes both uploading and downloading files from one host to another. FTP is both a protocol and an application. Specifically, FTP lets you copy files, list and manipulate directories, and view file contents. You can't use it to execute applications remotely.

Whenever a user attempts to access an FTP site, they will be asked to log in. If it's a public site, users can often just use the login name *anonymous* and then provide their email address as the password. Of course, there's no rule saying that they have to give their real email address if they don't want to. If the FTP site is secured, they will need a legitimate login name and password to access it. If they are using a browser such as Chrome, Firefox, or Edge to connect via FTP, the correct syntax in the address window is **`ftp://`*`username:password`*`@ftp.ftpsite.com`**.

The big downside to FTP is that it's unsecure. It transmits usernames and passwords in plain text. If a potential hacker is monitoring network traffic, this information will come through quite clearly. Be aware of this when using FTP, and make sure the FTP password is something not used to log into any other services. For secure file transfers, there are other options including Secure FTP (SFTP) and FTP Secure (FTPS).

In Windows, you can type a URL such as the one in the FTP example into the Search box to connect as well.

As mentioned earlier, FTP is both a protocol and an application, built into most operating systems. The built-in FTP app functionality is very basic—either the user uses a web browser or Windows File Explorer/macOS Finder to navigate the FTP server's folder structure, or they use the Command Prompt and text commands. Fortunately, there are also a number of third-party FTP client and server software packages available to make FTPing easier. Most of them are cross-platform (they work on multiple OSs), support security, have a graphical interface, allow for large file transfers, and allow for things like batch automation, searching, filtering, and bookmarking. We've listed a few of the most popular choices in Table 7.2. Figure 7.3 shows an example of the FileZilla interface.

TABLE 7.2 Popular FTP apps

Name	Website	OSs supported
FileZilla	`https://filezilla-project.org`	Windows, macOS, Linux, others
Cyberduck	`https://cyberduck.io`	Windows, macOS
WinSCP	`https://www.winscp.net`	Windows

Port 22—Secure Shell (SSH)

Secure Shell (SSH) is a connection-oriented protocol that can be used to set up a secure Telnet session for remote logins or for remotely executing programs and transferring files. Because it's secure, it was originally designed to be a replacement for the unsecure `telnet` command. A common client interface using SSH is called OpenSSH (`www.openssh.com`).

Port 23—Telnet

Speaking of *Telnet*, it seems that it has been around since the beginning of time as a terminal emulation protocol. Someone using Telnet can log into another machine and "see" the remote computer in a window on their screen. Although this vision is text only, the user can manage files on that remote machine just as if they were logged in locally.

FIGURE 7.3 FileZilla interface

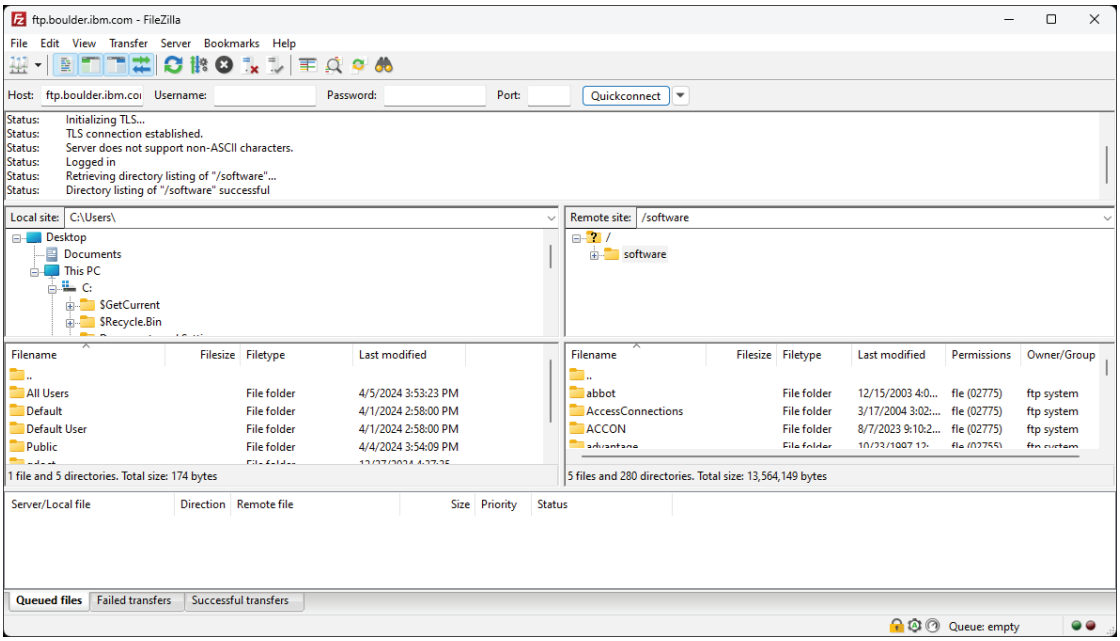
ftp.boulder.ibm.com - FileZilla
File Edit View Transfer Server Bookmarks Help
Host: ftp.boulder.ibm.co Username: Password: Port: Quickconnect
Status: Initializing TLS...
Status: TLS connection established.
Status: Server does not support non-ASCII characters.
Status: Logged in
Status: Retrieving directory listing of "/software"...
Status: Directory listing of "/software" successful
Local site: C:\Users\
Remote site: /software
Desktop
Documents
This PC
C:
$GetCurrent
$Recycle.Bin
software
Filename Filesize Filetype Last modified
All Users File folder 4/5/2024 3:53:23 PM
Default File folder 4/1/2024 2:58:00 PM
Default User File folder 4/1/2024 2:58:00 PM
Public File folder 4/4/2024 3:54:09 PM
1 file and 5 directories. Total size: 174 bytes
Filename Filesize Filetype Last modified Permissions Owner/Group
abbot File folder 12/15/2003 4:0... fle (02775) ftp system
AccessConnections File folder 3/17/2004 3:02:... fle (02775) ftp system
ACCON File folder 8/7/2023 9:10:2... fle (02775) ftp system
5 files and 280 directories. Total size: 13,564,149 bytes
Server/Local file Direction Remote file Size Priority Status
Queued files Failed transfers Successful transfers
Queue: empty

The problem with Telnet and other unsecure remote management options (such as RCP [remote copy] and FTP) is that the data they transmit, including passwords, is sent in plain text. Anyone eavesdropping on the line can intercept the packets and thus obtain usernames and passwords. SSH overcomes this by encrypting the traffic, including usernames and passwords.

Port 25—Simple Mail Transfer Protocol (SMTP)

This is the first of three protocols we'll look at devoted to email. *Simple Mail Transfer Protocol (SMTP)* is the protocol most commonly used to send email messages. Because it's designed to send only, it's referred to as a *push protocol*. SMTP is the protocol used to send email from mail server to mail server as well as from a mail server to an email client. An email client locates its email server by querying the DNS server for a mail exchange (MX) record. After the server is located, SMTP is used to push the message to the email server, which will then process the message for delivery.

We will cover DNS servers and MX records in more depth in Volume 1 Chapter 8, "Network Services."

Port 53—Domain Name System (DNS)

You probably use *Domain Name System (DNS)* every day whether you realize it or not. Its purpose is to resolve hostnames to IP addresses. For example, let's say that you open your web browser and type in a Uniform Resource Locator (URL) such as `www.wiley.com`. Your computer needs to know the IP address of the server that hosts that website in order for you

to connect to it. Through a DNS server, your computer resolves the URL to an IP address so communication can happen. DNS is so critical that we have an entire section dedicated to it in Volume 1 Chapter 8.

You will find some debate about what DNS stands for. The official name, according to the Internet Engineering Task Force (IETF), is Domain Name System. However, you may see others reference it as Domain Name Server as well. Regardless, the important thing to remember is what DNS does, which is resolving hostnames to IP addresses.

Port 67/68—Dynamic Host Configuration Protocol (DHCP)

Dynamic Host Configuration Protocol (DHCP) dynamically assigns IP addresses and other IP configuration information to network clients. Configuring your network clients to receive their IP addresses from a DHCP server reduces network administration headaches. We'll cover the mechanics of how DHCP works in Volume 1 Chapter 8, along with DNS.

Port 80—Hypertext Transfer Protocol (HTTP)

HTTP was once the most used Process/Application layer protocol. It manages the communication between a web server and client, and it lets you connect to and view all the content that you enjoy on the Internet. All the information transmitted by HTTP is plain text, which means that it's not secure. Therefore, it's not a good choice for transmitting sensitive or personal information, such as usernames and passwords, or for transmitting banking information. Because of that, it's been supplanted by HTTPS (covered later).

Port 110—POP3

For a long time, *Post Office Protocol 3 (POP3)* was the preferred protocol for downloading email. It's been replaced in most installations by IMAP4 (covered later) because IMAP4 includes security and more features than POP3.

Port 137-139—Network Basic Input/Output System (NetBIOS)/NetBIOS over TCP/IP (NetBT)

Network Basic Input/Output System (NetBIOS) is an application programming interface (API) that allows computers to communicate with each other over the network. It works at Layer 5 of the OSI model. Consequently, it needs to work with another network protocol to handle the functions of Layer 4 and below. NetBIOS running over TCP/IP is called *NetBT*, or *NBT*. Specifically, NetBIOS provides three services:

- Naming service, for name registration and resolution
- Datagram distribution service, for connectionless communication
- Session management service, for connection-oriented communication

For many years, Microsoft network clients were configured with a NetBIOS name, which was their network name. To communicate with another computer on the network, the

NetBIOS name would need to be resolved (matched) to an IP address. This was done with a WINS (Windows Internet Name Service) server or `LMHOSTS` file and could not be performed across any routed connection (which includes the Internet).

If you'll recall, back in the "Port 53—DNS" section we introduced the concept of a hostname, which is the computer's name on a network. The NetBIOS name is analogous and could be one and the same as the hostname or something totally different. (Two different computer names for the same computer on a network? What could go wrong?) The big differences are that hostnames are resolved with a DNS server (or `HOSTS` file) and can work across the Internet. WINS was far inferior to DNS for name resolution, so Microsoft ended up adopting hostnames and DNS like the rest of the industry. We will also talk about hostnames again in Volume 2 Chapter 3, "Windows Administration."

Port 143—Internet Message Access Protocol (IMAP)

Internet Message Access Protocol (IMAP) is a secure protocol designed to download email. Its current version is version 4, or IMAP4. It's the client-side email management protocol of choice, having replaced the unsecure POP3. Most current email clients, such as Microsoft Outlook and Gmail, are configured to be able to use either IMAP4 or POP3. IMAP4 has some definite advantages over POP3. They include the following:

- IMAP4 works in connected and disconnected modes. With POP3, the client makes a connection to the email server, downloads the email, and then terminates the connection. IMAP4 allows the client to remain connected to the email server after the download, meaning that as soon as another email enters the inbox, IMAP4 notifies the email client, which can then download it.
- IMAP4 also lets you store the email on the server, as opposed to POP3, which requires you to download it.
- IMAP4 allows multiple clients to be simultaneously connected to the same inbox. This can be useful for smartphone users who have both Outlook on their workstation and their smartphone email client operational at the same time or for cases where multiple users monitor the same mailbox, such as on a customer service account. IMAP4 allows each connected user or client to see changes made to messages on the server in real time.

The A+ exam objectives incorrectly spell out IMAP as Internet Mail Access Protocol. As we've stated previously, focus on what the acronym "does" instead of what it specifically stands for and you'll be in good shape.

Port 161/162—Simple Network Management Protocol (SNMP)

We said you needed to memorize all of the protocols and ports in these sections, but here we include one that's no longer listed in exam objectives. It's still commonly used for network monitoring and management though, so we think it's important for you to be aware of it. It's called *Simple Network Management Protocol (SNMP)*, and it gathers and manages network performance information.

On your network, you might have several connectivity devices, such as routers and switches. A management device called an *SNMP server* can be set up to collect data from

these devices (called *agents*) and ensure that your network is operating properly. Although SNMP is mostly used to monitor connectivity devices, many other network devices are SNMP-compatible as well. The most current version is SNMPv3.

SNMPv3 allows for the use of strong cryptographic algorithms such as Advanced Encryption Standard (AES), hash-based message authentication code - secure hash algorithm (HMAC-SHA), and hash-based message authentication code - message-digest algorithm 5 (HMAC-MD5). This helps ensure email data integrity and confidentiality.

Port 389—Lightweight Directory Access Protocol (LDAP)

The *Lightweight Directory Access Protocol (LDAP)* is a directory services protocol based on the X.500 standard. LDAP is designed to access information stored in an information directory typically known as an LDAP directory or LDAP database.

X.500 is an industrywide standard for organizing and storing information about users, services, and resources on a network.

On your network, you probably have a lot of information, such as employee phone numbers and email addresses, client contact lists, and infrastructure and configuration data for the network and network applications. This information might not get updated frequently, but you might need to access it from anywhere on the network, or you might have a network application that needs access to this data. LDAP provides you with the access, regardless of the client platform from which you're working. You can also use access control lists (ACLs) to set up who can read and change entries in the database using LDAP. A common analogy is that LDAP provides access to and the structure behind your network's phone book.

Port 443—Hypertext Transfer Protocol Secure (HTTPS)

To encrypt traffic between a web server and client securely, *Hypertext Transfer Protocol Secure (HTTPS)* can be used. HTTPS connections are secured using either *Secure Sockets Layer (SSL)* or *Transport Layer Security (TLS)*. SSL and TLS have the job of encrypting data in transit between the server and client.

Real World Scenario

How Secure Is It?

You have probably heard before that you should not enter personal information (such as a credit card number) into an unsecure website. But what does that really mean?

First, know what to look for. If you were to enter information into a website with an address beginning with just `http://`, you would be asking for someone to steal the information! The HTTP protocol transmits data in plain text, meaning that there is no encryption at all

between your computer and the server. On the other hand, HTTPS encrypts the data transmissions as they cross the wire (known as *data in transit*).

To use HTTPS, the website needs to obtain an SSL/TLS certificate from a reputable, trusted organization called a certificate authority (CA), which verifies the identity of the website. So the good news is that if you are accessing a site with `https://` in the header, you know that the site is what it says it is (and not a fake site or Trojan horse) and that transmissions between your computer and that site are encrypted.

So to answer the question posed by this sidebar, if a site uses HTTPS, you can be assured that the data in transit is encrypted (and secure). If you trust the receiving website, then it's totally fine to enter sensitive information. With the volume of Internet commerce (or e-commerce) in the world today, this protocol is critical. Keep in mind, though, that once the data is on the website's server—that is, the data is "at rest"—HTTPS is no longer relevant and other protection methods are needed to keep your data secure. The data breaches you hear about in the news almost always target data at rest and not data in transit.

Occasionally, you might visit a website that uses HTTPS and get a pop-up error message saying that the certificate has expired or could not be validated. This is most likely a case of the certificate legitimately expiring, but it could be that it's a Trojan horse website. Proceed with caution!

Port 445—Server Message Block (SMB)/Common Internet File System (CIFS)

Server Message Block (SMB) is a protocol originally developed by IBM but then enhanced by Microsoft, IBM, Intel, and others. It's used to provide shared access to files, printers, and other network resources and is primarily implemented by Microsoft systems. In a way, it can function a bit like FTP only with a few more options, such as the ability to connect to printers, and more management commands. It's also known for its ability to make network resources easily visible through various Windows network apps (such as Network in File Explorer).

Common Internet File System (CIFS) is a Microsoft-developed enhancement of the SMB protocol. The intent behind CIFS is that it can be used to share files and printers between computers, regardless of the operating system that they run. It's the default file and print sharing protocol in Windows.

Port 3389—Remote Desktop Protocol (RDP)

Developed by Microsoft, the *Remote Desktop Protocol (RDP)* allows users to connect to remote computers and run programs on them. When you use RDP, you see the desktop of

the computer you've signed into on your screen. It's like you're really there, even though you're not.

When you use RDP, the computer at which you are seated is the client, and the computer you're logging into is the server. RDP client software is available for Windows, Linux, macOS, iOS, and Android. Microsoft's RDP client software is called *Remote Desktop Connection*. The server uses its own video driver to create video output and sends the output to the client using RDP. Conversely, all keyboard and mouse input from the client is encrypted and sent to the server for processing. RDP also supports sound, drive, port, and network printer redirection. In a nutshell, this means that if you could see, hear, or do it sitting at the remote computer, you could see, hear, or do it at the RDP client too.

Services using this protocol can be great for telecommuters. It's also very handy for technical support folks, who can log into and assume control over a remote computer. It's a lot easier to troubleshoot and fix problems when you can see what's going on and "drive."

Understanding IP Addressing

To communicate on a TCP/IP network, each device needs to have a unique IP address. Any device with an IP address is referred to as a *host*. This can include servers, workstations, printers, routers, and other devices. If you can assign it an IP address, it's a host. As an administrator, you can assign the host's IP configuration information manually (a *static IP address*), or you can have it automatically assigned by a DHCP server (a *dynamic IP address*). On the client, this is done through the Network & Internet app in Settings. Exercise 7.1 shows you how to change the client computer's IP address assignment.

Exam Objective 2.6, "Given a scenario, configure basic wired/wireless small office/home office (SOHO) networks," states that you should know the difference between static (manually configured) and dynamic (DHCP supplied) IP addresses.

EXERCISE 7.1

Changing IP Address Assignment in Windows 11

The IP address on a computer can be assigned manually by an administrator or automatically if the network has a DHCP server. This exercise shows you how to check (and change) IP configuration information.

1. Open the Settings app by clicking Start, then Settings.
2. Select Network & Internet from the menu on the left side. You will get a screen like the one shown in Figure 7.4.

EXERCISE 7.1 *(continued)*

FIGURE 7.4 Network & Internet

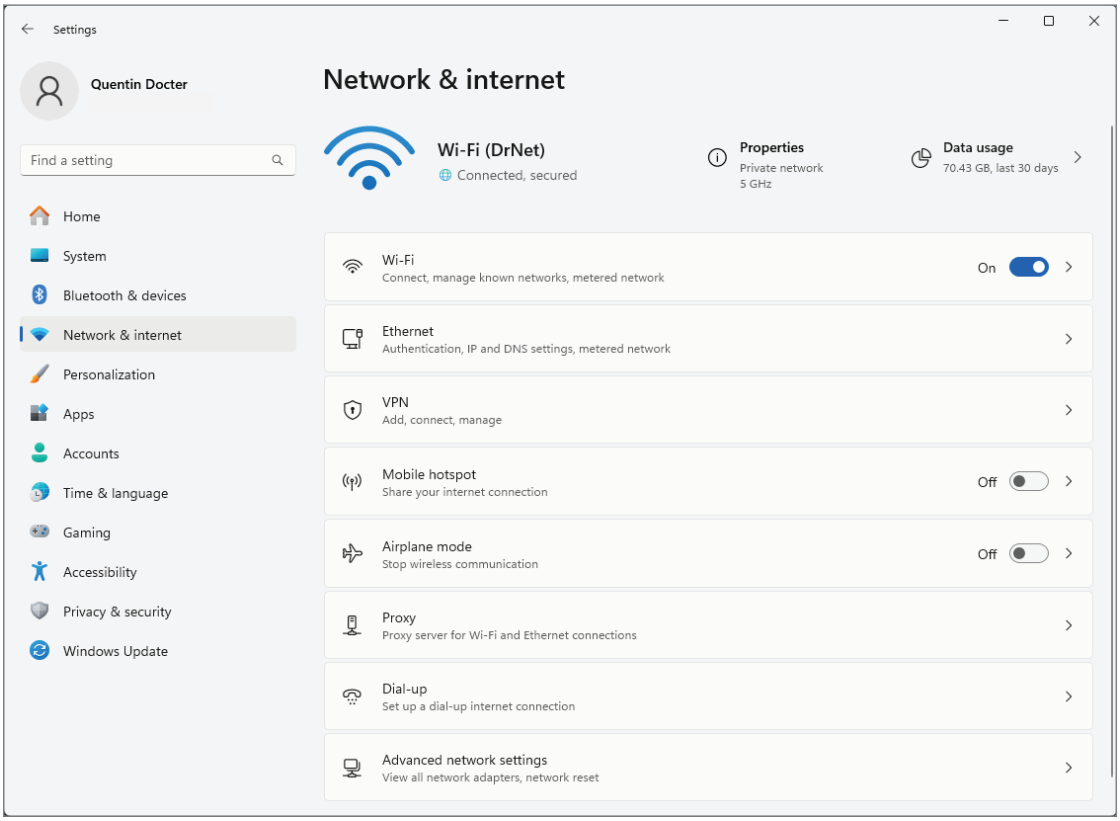

Using Figure 7.4 as a reference, you have a few different ways to get to the IP settings. One way is to click on either Wi-Fi or Ethernet at the top of the list, depending on the connection you want to manage. The other is to use the Advanced Network Settings option at the bottom. Any of those will eventually get you to the same place.

3. Click Advanced Network Settings. Here you will see the list of available network adapters in the Network Adapters section at the top.

4. To manage the Wi-Fi connection, click the down arrow on the right side of the Wi-Fi box (Figure 7.5).

FIGURE 7.5 Advanced Network Settings

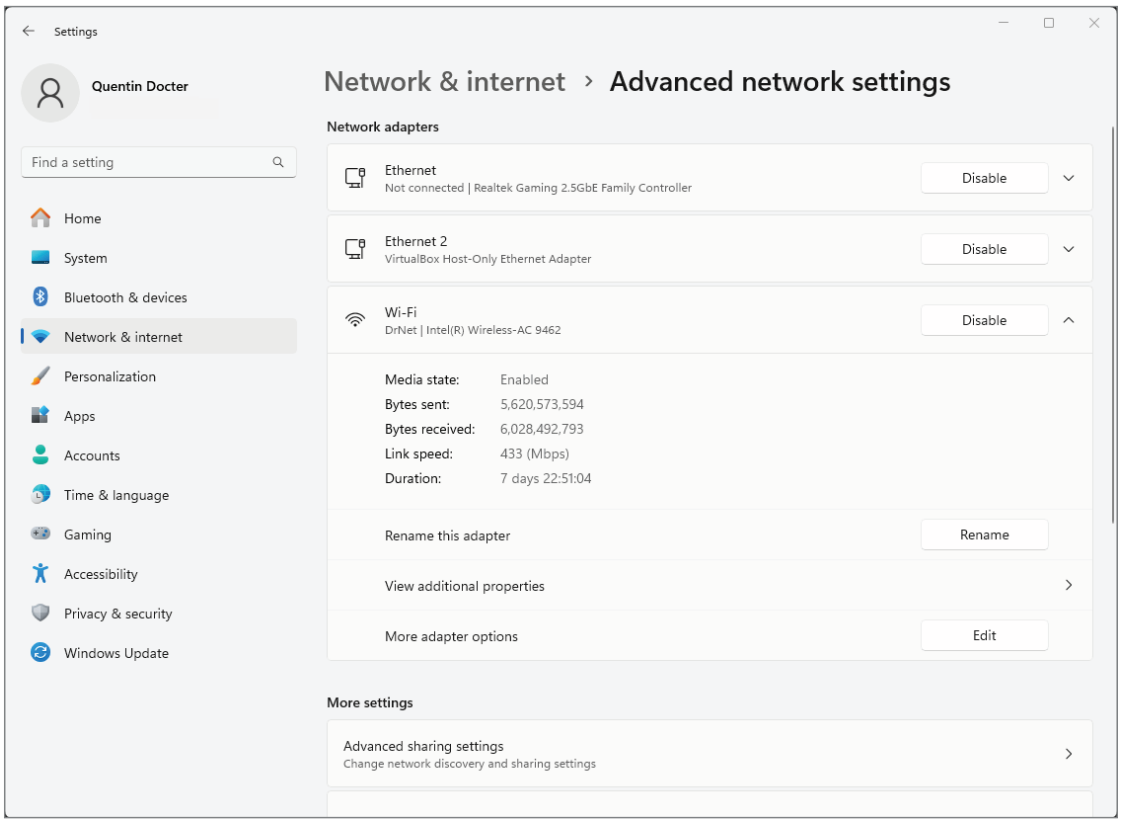

5. Click View Additional Properties to see the properties of this connection (Figure 7.6).

EXERCISE 7.1 *(continued)*

FIGURE 7.6 Wi-Fi connection properties

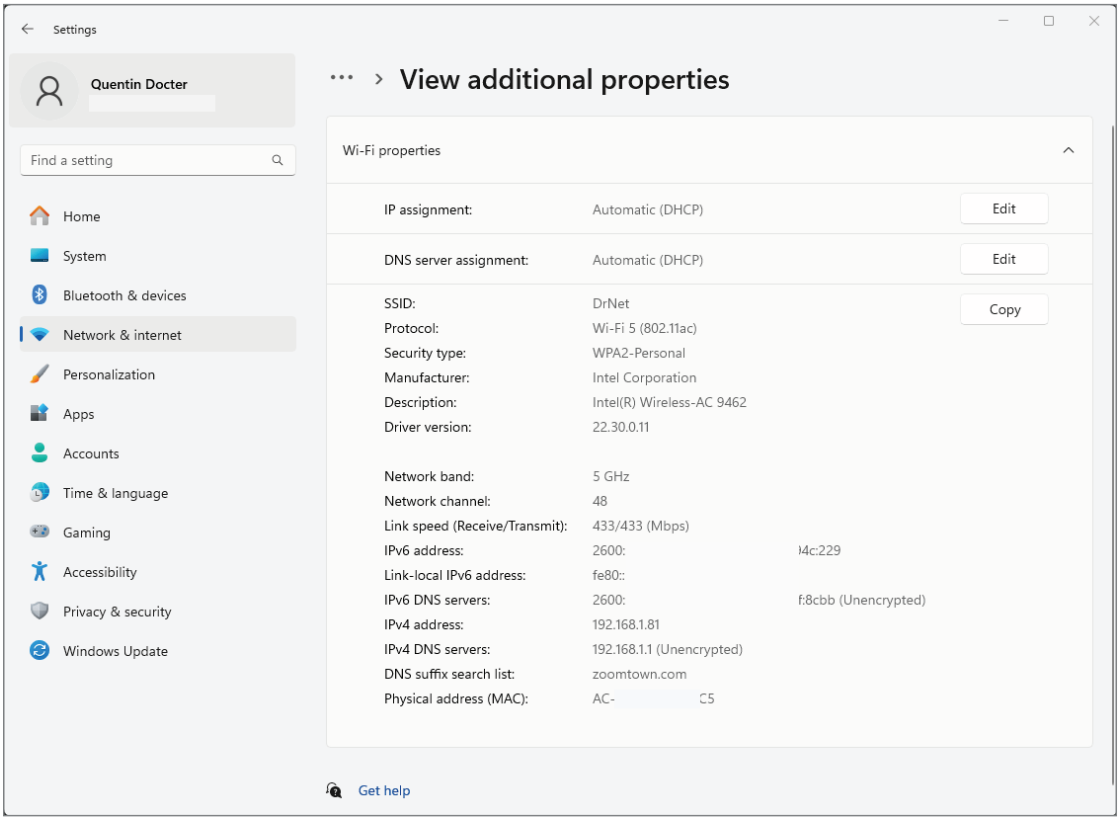

In Figure 7.6, notice that the IP assignment is set to Automatic (DHCP). This is the preferred method if the network has a DHCP server.

6. To change the setting, click the Edit button to the right of the IP assignment.
7. In the Edit IP Settings box, change the option to Manual. The open dialog window will automatically change to the next choice.
8. In the second Edit IP Settings box, slide the toggle on for IPv4. You will see a screen like the one in Figure 7.7.

FIGURE 7.7 Manually configuring IP settings

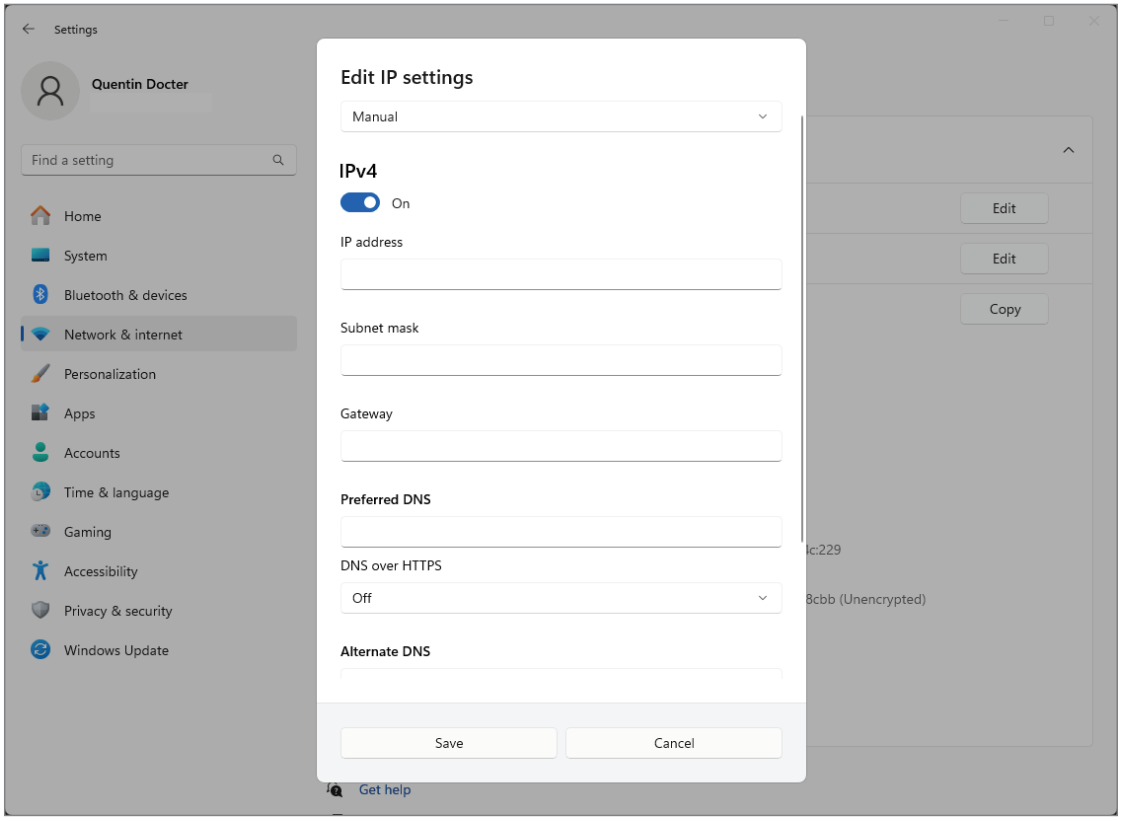

9. Enter the IP address, subnet mask, gateway, and preferred DNS information.

   The IP address and subnet mask are required to participate on a network. To communicate with computers on another network (on the other side of a router), the gateway is also needed. (We will explain what subnet masks and gateways are shortly.) DNS servers are incredibly useful for resolving host names to IP addresses, so you don't have to memorize a long list of IP addresses to access resources.

10. When finished, click Save.

    If you are working with a version of Windows older than Windows 10, IP configuration information was managed through the network adapter's IP Properties section. To get to a similar screen as to what you'd see (but in Windows 11), refer to Figure 7.5 and look for More Adapter Options. Click the Edit button on the right to get the Properties window shown in Figure 7.8.

EXERCISE 7.1 *(continued)*

FIGURE 7.8 Network connection properties

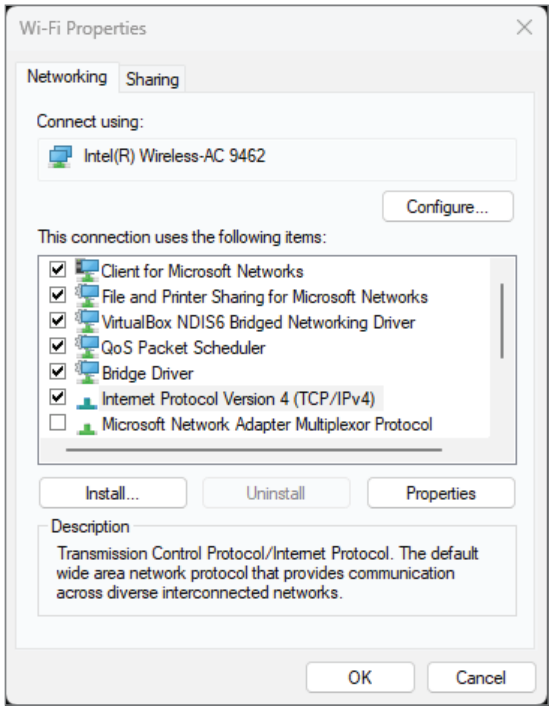

Highlight Internet Protocol Version 4 (TCP/IPv4) and click Properties again (Figure 7.9).

FIGURE 7.9 IPv4 properties

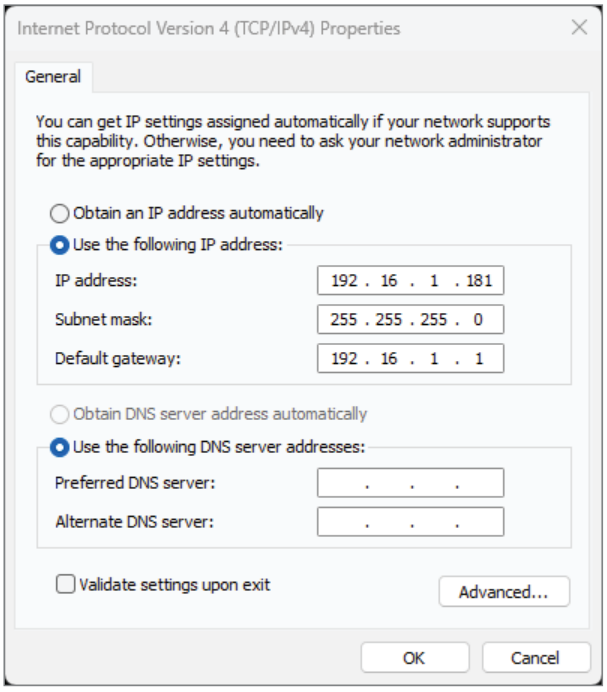
Internet Protocol Version 4 (TCP/IPv4) Properties
General
You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.
Obtain an IP address automatically
Use the following IP address:
IP address: 192 . 16 . 1 . 181
Subnet mask: 255 . 255 . 255 . 0
Default gateway: 192 . 16 . 1 . 1
Obtain DNS server address automatically
Use the following DNS server addresses:
Preferred DNS server:
Alternate DNS server:
Validate settings upon exit
Advanced...
OK
Cancel

For manual configuration, click Use The Following IP Address and then enter the required configuration information. DHCP clients should use Obtain An IP Address Automatically. Click OK when finished.

We also mentioned earlier that you can get to the Windows 11 IP configuration options by choosing the network connection from the Network & Internet screen shown in Figure 7.4. To do so from there, click the name of the connection (such as Wi-Fi), then click Hardware Properties. The screen you will see is nearly identical to the one shown in Figure 7.6. From there, follow steps 6 through 10 like you did in this exercise to change the settings.

An IPv4 address is a 32-bit hierarchical address that identifies a host on the network. It's typically written in dotted-decimal notation, such as 192.168.10.55. Each of the numbers in this example represents 8 bits (or 1 byte) of the address, also known as an *octet*. The same address written in binary (how the computer thinks about it) would be 11000000 10101000 00001010 00110111. As you can see, the dotted-decimal version is a much more convenient way to write these numbers.

This section covers IPv4. IPv6 will be covered in its own section.

The addresses are said to be hierarchical, as opposed to "flat," since the numbers at the beginning of the address identify groups of computers that belong to the same network. Because of the hierarchical address structure, we're able to do really cool things, such as route packets between local networks and on the Internet.

A great example of hierarchical addressing is your street address. Let's say that you live in apartment 4B at 123 Main Street, Anytown, Kansas, USA. If someone sent you a letter via snail mail, the hierarchy of your address helps the postal service and carrier deliver it to the right place. First and broadest is USA. Kansas helps narrow it down a bit, and Anytown narrows it down more. Eventually we get to your street, the right number on your street, and then the right apartment. If the address space were flat (for example, Kansas didn't mean anything more specific than Main Street), or you could use any name you wanted for your state, it would be really hard to get the letter to the right spot.

Take this analogy back to IP addresses. They're set up to organize networks logically in order to make delivery between them possible and then to identify an individual node within a network. If this structure weren't in place, a huge, multinetwork space like the Internet probably wouldn't be possible. It would simply be too unwieldy to manage.

Another example of a hierarchical addressing scheme is telephone numbers. The first three digits, the area code, group all telephone numbers with that area code into one logical network. The second grouping of three numbers defines a local calling area, and the last grouping of numbers is the unique identifier within that local calling area.

A Quick Binary Tutorial

As we mentioned earlier, each IP address is written in four octets in dotted-decimal notation, but each octet represents 8 bits. A binary bit is a value with two possible states: on equals 1 and off equals 0. If the bit is turned on, it has a decimal value based on its position within the octet. An off bit always equals 0. Take a look at Figure 7.10, which will help illustrate what we mean.

If all the bits in an octet are off, or 00000000, the corresponding decimal value is 0. If all bits in an octet are on, you would have 11111111, which is 255 in decimal.

When you're working with IPv4 addressing, all numbers will be between 0 and 255.

FIGURE 7.10 Binary values

Position in octet	8	7	6	5	4	3	2	1
Bit on	1	1	1	1	1	1	1	1
Has the decimal value of . . .	128	64	32	16	8	4	2	1
Mathematically	2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0

Where it starts to get more entertaining is when you have combinations of zeroes and ones. For example, 10000001 is equal to 129 (128 + 1), and 00101010 is equal to 42 (32 + 8 + 2).

As you work with IPv4 addresses, you'll see certain patterns emerge. For example, you may be able to count quickly from left to right in an octet pattern, such as 128, 192, 224, 240, 248, 252, 254, and 255. That's what you get if you have (starting from the left) 1, 2, 3, and so forth up to 8 bits on in sequence in the octet. In fact, those numbers are good to commit to memory if you're dealing with IPv4 a lot, or about to take an exam that will test you on it. You will see why, a few pages later in this chapter in Table 7.4.

It's beyond the scope of this book to get into too much detail on binary-to-decimal conversion, but this primer should get you started.

If you want more information on binary, here are two helpful resources.

A network binary math lesson can be found at `https://learningnetwork.cisco.com/s/blogs/a0D3i000002SKMwEAO/network-binary-math-explained`, and you'll find a binary-to-decimal converter at `https://www.rapidtables.com/convert/number/decimal-to-binary.html`.

Parts of the IP Address

Each IP address is made up of two components: the *network ID* and the *host ID*. The network portion of the address always comes before the host portion. Because of the way IP addresses are structured, the network portion does not have to be a specific fixed length. In other words, some computers will use 8 of the 32 bits for the network portion and the other 24 for the host portion, whereas other computers might use 24 bits for the network portion and the remaining 8 bits for the host portion. Here are a few rules that you should know about when working with IP addresses:

- All host addresses on a network must be unique.
- On a routed network (such as the Internet), all network addresses must be unique as well.
- Neither the network ID nor the host ID can be set to all 0s. A host ID portion of all 0s means "this network."
- Neither the network ID nor the host ID can be set to all 1s. A host ID portion of all 1s means "all hosts on this network," commonly known as a *broadcast address*.

The term "broadcast" on a network generally has very bad connotations. It means that every computer on that network segment must accept and receive the message whether it's useful or not. It's kind of like email spam or junk postal mail that you have to look at and then throw away. There are some instances where broadcasts are useful (e.g., when a client is looking for a DHCP server), but in general they should be minimized to reduce network traffic.

Computers are able to differentiate where the network ID ends and the host address begins through the use of a *subnet mask*. This is a value written just like an IP address and may look something like 255.255.255.0. Any bit that is set to a 1 in the subnet mask makes the corresponding bit in the IP address part of the network ID (regardless of whether the bit in the IP address is on or off). When setting bits to 1 in a subnet mask, you always have to turn them on sequentially from left to right, so that the bits representing the network address are always contiguous and come first. The rest of the address will be the host ID. The number 255 is the highest number you will ever see in IP addressing, and it means that all bits in the octet are set to 1.

Here's an example based on two numbers that we have used in this chapter. Look at the IP address of 192.168.10.55. Let's assume that the subnet mask in use with this address is 255.255.255.0. This indicates that the first three octets are the network portion of the address and the last octet is the host portion; therefore, the network portion of this ID is 192.168.10 and the host portion is 55. If the subnet mask were 255.255.0.0, the computer would see its network address as 192.168 and its host address as 10.55. As you can see, the subnet mask can make the exact same address appear as though it's on a different network.

If you're ever dealing with network communication issues, the IP address and subnet mask are among the first things you should check.

To communicate using IPv4, each computer is *required* to have an IP address and correct subnet mask. A third component, called a *default gateway*, identifies the IP address of the device that will allow the host to connect outside of the local network. This is typically your router, and it's required if you want to communicate with computers outside of your local network. An example configuration was shown back in Figure 7.9.

For the exam, remember that Exam Objective 2.6 specifies that you should know what subnet masks and gateways are, and how to configure them on a basic SOHO network.

IPv4 Address Classes

The designers of TCP/IP designated classes of networks based on the first 3 bits of the IP address. As you will see, classes differ in how many networks of each class can exist and the number of unique hosts that each network can accommodate. Here are some characteristics of the three classes of addresses that you will commonly deal with:

Class A Class A networks are defined as those with the first bit set as 0 (decimal values from 0 to 127) and are designed for very large networks. The default network portion for Class A networks is the first 8 bits, leaving 24 bits for host identification. Because the network portion is only 8 bits long (and 0 and 127 are reserved), there are only 126 Class A network addresses available. The remaining 24 bits of the address allow each Class A network to hold as many as 16,777,214 hosts. Examples of Class A networks include the networks for telecommunications giants Comcast and AT&T and organizations such as `Amazon.com`, IBM, Hewlett-Packard, Apple, Ford, and the U.S. Department of Defense. Most Class A blocks are assigned to regional Internet registries, who allocate IP addresses to a specific region of the world. All possible Class A networks are in use; no more are available.

The number of networks available is determined by the formula 2^n, where n represents the number of bits being used. In the Class A example, 7 bits are available by default (because the first one is always set as 0 by design), so there are 2^7 networks available, which is 128. However, the network addresses of 0 and 127 are also reserved, so it's really 126.

The number of hosts available is determined by the formula $2^n - 2$, because a host address of all 0s or all 1s is not allowed. Remember, all 0s means "this network" and all 1s are broadcast addresses. So, in a default Class A network, there can be $2^{24} - 2$ hosts, or 16,777,214.

Class B Class B networks always have the first 2 bits set at 10 (decimal values from 128 to 191) and are designed for medium-sized networks. The default network portion for Class B networks is the first 16 bits, leaving 16 bits for host identification. This allows for 16,384 (2^{14}) networks, each with as many as 65,534 ($2^{16} - 2$) hosts attached. Examples of Class B networks include the networks of Microsoft, ExxonMobil, and Purdue University. Class B networks are generally regarded as unavailable, but address-conservation techniques have made some of these addresses available from time to time over the years.

Class C Class C networks have the first three bits set at 110 (decimal values from 192 to 223) and are designed for smaller networks. The default network portion for Class C networks is the first 24 bits, leaving 8 bits for host identification. This allows for 2,097,152 (2^{21}) networks, but each network can have a maximum of only 254 ($2^8 - 2$) hosts. Most companies have Class C network addresses. A few class C networks are still available.

The address assignment examples in this chapter refer to addresses that are used on the Internet. For example, Apple has the network address of 17.0.0.0. No one else on the Internet can use addresses in that network's range (17.0.0.0–17.255.255.255). But if you are using IP addresses on an internal network that never connects to the Internet, you are free to use whatever addresses you would like.

Table 7.3 shows the IPv4 classes, their ranges, and their default subnet masks.

The network addresses 0 and 127 are reserved and not available for use. Specifically, the address 127.0.0.1, called the *loopback address*, is used for troubleshooting network adapters. We'll talk more about this in Volume 1 Chapter 12, "Core Hardware and Network Troubleshooting."

TABLE 7.3 IPv4 address classes

Class	First octet	Default subnet mask	Comments
A	1–127	255.0.0.0	For very large networks; 127 is reserved for the loopback address.
B	128–191	255.255.0.0	For medium-sized networks.
C	192–223	255.255.255.0	For smaller networks with fewer hosts.
D	224–239	N/A	Reserved for multicasts (sending messages to multiple systems).
E	240–255	N/A	Reserved for testing.

The IP address is often written in shorthand to show how many bits are being used for the network portion of the address. For example, you might see something like 10.0.0.0/8. The /8 on the end indicates that the first 8 bits are the network portion of the address, and the other 24 are the host portion. Another example is 192.168.1.0/24, which is a Class C network with a default subnet mask.

Classless Inter-Domain Routing

The default subnet masks for each class of address are by no means the only subnet masks that can be used. In fact, if they were, it would severely limit the number of possible TCP/IP networks available. To resolve this and provide additional addressing flexibility, there is *classless inter-domain routing (CIDR)*. This is just a fancy of way of saying, "You don't have to use the default subnet masks." From a practical standpoint, CIDR minimizes the concept of IP address classes and primarily focuses on the number of bits that are used as part of the network address.

Taking a look at the defaults can help illustrate how CIDR works. If you have a Class A default mask of 255.0.0.0, that is 11111111.00000000.00000000.00000000 in binary. A Class B default mask of 255.255.0.0 is 11111111.11111111.00000000.00000000 in binary. There's no rule that says you have to use an entire octet of bits to represent the network portion of the address. The only rule is that you have to add 1s in a subnet mask from left to right. What if you wanted to have a mask of 255.240.0.0 (11111111.11110000.00000000 .00000000); can you do that? The answer is yes, and that is essentially what CIDR does. Table 7.4 shows you every available subnet mask and its equivalent slash notation.

TABLE 7.4 CIDR values

Subnet mask	Notation
255.0.0.0	/8
255.128.0.0	/9
255.192.0.0	/10
255.224.0.0	/11
255.240.0.0	/12
255.248.0.0	/13
255.252.0.0	/14
255.254.0.0	/15
255.255.0.0	/16

Subnet mask	Notation
255.255.128.0	/17
255.255.192.0	/18
255.255.224.0	/19
255.255.240.0	/20
255.255.248.0	/21
255.255.252.0	/22
255.255.254.0	/23
255.255.255.0	/24
255.255.255.128	/25
255.255.255.192	/26
255.255.255.224	/27
255.255.255.240	/28
255.255.255.248	/29
255.255.255.252	/30

Earlier, we said that CIDR minimizes the impact of classes, but there are still some restrictions. The /8 through /15 notations can be used only with Class A network addresses; /16 through /23 can be used with Class A and B network addresses; /24 through /30 can be used with Class A, B, and C network addresses. You can't use anything more than /30, because you always need at least 2 bits for hosts.

Now that you know that you *can* do it, the question is, *why* would you do it? The answer is that it provides you with the flexibility to configure your network.

Here's an example. Say that your default network address is 10.0.0.0/8. That means that you have 24 bits left for hosts on that one network, so you can have just over 16.7 million hosts. How realistic is it that one company will have that many hosts? It's not realistic at all, and that doesn't even bring up the issue that the network infrastructure wouldn't be able to handle physically having that many hosts on one network. However, let's say that you work for a large corporation with about 15 divisions and some of them have up to 3,000 hosts. That's plausible. What you can do is to set up your network so that each division has its own

smaller portion of the network (a subnet) big enough for its needs. To hold 3,000 hosts and have a bit of room for expansion, you need 12 bits ($2^{12} - 2 = 4{,}094$), meaning that you have 20 bits left over for the network address. Thus, your new configuration could be 10.0.0.0/20.

Public vs. Private IP Addresses

All the addresses that are used on the Internet are called *public addresses*. They must be purchased, and only one computer can use any given public address at one time. The problem that presented itself was that the world was soon to run out of public IP addresses while the use of TCP/IP was growing. And that was back in the 20th century, when smart devices and IoT didn't exist—they need about 30 billion IP addresses (as of 2025, and the number is growing dramatically) themselves. Additionally, the structure of IP addressing made it impossible to "create" or add any new addresses to the system.

Exam Objective 2.6 calls out public addresses and private addresses as topics to know for the test. Understand how they are different and where each is used.

To address this, a solution was devised to allow for the use of TCP/IP without requiring the assignment of a public address. The solution was to use private addresses. *Private addresses* are not routable on the Internet. They were intended for use on private networks only. That private addresses weren't intended for use on the Internet freed us from the requirement that all addresses be globally unique. This essentially created an infinite number of IP addresses that companies could use within their own network walls.

Although this solution helped alleviate the problem of running out of addresses, it created a new one. The private addresses that all of these computers have aren't globally unique, but they need to be in order to access the Internet.

A service called *Network Address Translation (NAT)* was created to solve this problem. NAT runs on your router and handles the translation of private, nonroutable IP addresses into public IP addresses. There are three ranges reserved for private, nonroutable IP addresses, as shown in Table 7.5. You should memorize these ranges and be able to identify them on sight.

TABLE 7.5 Private IP address ranges

Class	IP address range	Default subnet mask	Number of hosts
A	10.0.0.0–10.255.255.255	255.0.0.0	16.7 million
B	172.16.0.0–172.31.255.255	255.240.0.0	1 million
C	192.168.0.0–192.168.255.255	255.255.0.0	65,534

Real World Scenario

Private IP Addresses and Subnet Masks

When you look at the default subnet masks for the private IP address ranges, you might think, "Wait a minute. Those masks aren't the same as the default subnet masks for the address class," and you'd be correct.

To understand how TCP/IP addresses work, it's often helpful to start with the concept of address classes, because it helps you break the information into chunks, making it easier to understand. In the real world, though, most network administrators don't think in terms of classes, and routers certainly don't operate based on classes. Communication and routing on a network all happen in binary. Experienced network admins will think in terms like, "I am dealing with a 10.0.0.0/16 network." They know the address and the length of the subnet mask.

Earlier in this chapter, you learned about the concept of CIDR, which basically ignores the artificial boundaries of address classes. It uses a concept called *variable length subnet masking (VLSM)*, which might sound complicated, but it just means that the length of the subnet mask determines the structure of the network. (And by structure, we mean the network addresses and the number of networks and hosts that you can have on a network.)

How does this relate back to private IP address ranges? You'll notice that the Class A address range is 10.0.0.0/8, which has a "default" mask for a Class A address. 172.16.0.0/12 is an address in the Class B range, but it does not use the "default" /16 mask. If it did use a /16 mask, then the administrator would have only the remaining 16 bits to use for additional subnets and hosts. As it is, the administrator has 20 bits to play with, which provides much greater flexibility in designing the network. The same concept applies to 192.168.0.0/16. The administrator has 16 free bits to create subnets and host ranges, whereas the "default" /24 mask would leave only 8 bits and not a lot of flexibility.

There are three things that you should take away from this sidebar:

- Know the subnet masks in Table 7.5 and understand that they are different from the default masks for that class of address.
- Know that you are not limited to using the default masks or class rules.
- It's all about the binary.

The A+ exam may test you on the basics of IP addressing and subnetting, which we have covered in this book. If you pursue more advanced certifications, such as the CompTIA Network+ or the Cisco series of certifications, you will be expected to know IP addressing and subnetting in depth. If you are interested in learning more (after you pass the A+ exam, of course), check out *CompTIA Network+ Study Guide*, by Todd Lammle and Jon Buhagiar (May 7, 2024, Sybex).

These private addresses cannot be used on the Internet and cannot be routed externally. The fact that they are not routable on the Internet is actually an advantage because a network administrator can use them essentially to hide an entire network from the Internet.

This is how it works: The network administrator sets up a NAT-enabled router, which functions as the default gateway to the Internet. The external interface of the router has a public IP address assigned to it that has been provided by the ISP, such as 155.120.100.1. The internal interface of the router will have an administrator-assigned IP address within one of the private ranges, such as 192.168.1.1. All computers on the internal network will then also need to be on the 192.168.1.0 network. To the outside world, any request coming from the internal network will appear to come from 155.120.100.1. The NAT router translates all incoming packets and sends them to the appropriate client. This type of setup is very common today.

By definition, NAT is actually a one-to-one private-to-public IP address translation protocol. There is a type of NAT called *NAT Overload*, also known as *Port Address Translation (PAT)*, which allows for many private IP addresses to use one public IP address on the Internet. You probably don't need to have this depth of knowledge for the A+ exam, but it's good stuff to be aware of.

You may look at your own computer, which has an address in a private range, and wonder, "If it's not routable on the Internet, then how am I on the Internet?" Remember, the NAT router technically makes the Internet request on your computer's behalf, and the NAT router is using a public IP address.

Don't make the mistake of thinking that your internal network can't be hacked if it is using private addresses through NAT. It can. Hackers just have to use more tools and try a little harder to uncover your internal structure. Even if you're using NAT, you still need protective features such as firewalls and antimalware software.

Automatic Private IP Addressing

Automatic Private IP Addressing (APIPA) is a TCP/IP standard used to automatically configure IP-based hosts that are unable to reach a DHCP server. APIPA addresses are in the 169.254.0.0–169.254.255.255 range, with a subnet mask of 255.255.0.0. If you see a computer that has an IP address beginning with 169.254, you know that it has configured itself.

APIPA is a specific objective called out in exam objective 2.6. Be sure to know how it works.

Typically, the only time that you will see this is when a computer is supposed to receive configuration information from a DHCP server but for some reason that server is unavailable. Even while configured with this address, the client will continue to broadcast for a DHCP server so that it can be given a real address once the server becomes available.

APIPA is also sometimes known as *zero configuration networking* or *address autoconfiguration*. Both of these terms are marketing efforts, created to remove the perceived difficulty of configuring a TCP/IP network. While TCP/IP has generally been considered difficult to configure (compared to other protocols), APIPA can make it so that a TCP/IP network can run with no configuration at all! For example, say that you are setting up a small local area network that has no need to communicate with any networks outside of itself. To accomplish this, you can use APIPA to your advantage. Set the client computers to receive DHCP addresses automatically, but don't set up a DHCP server. The clients will configure themselves and be able to communicate with each other using TCP/IP. The only downside is that this will create a little more broadcast traffic on your network. This solution is only really effective for a nonrouted network of fewer than 100 computers. Considering that most networks today need Internet access, it's unlikely that you'll run across a network configuration like this.

Real World Scenario

Help! I Can't Get to the Internet!

This is something that you will probably hear a lot: A user on your network calls and complains that they can't get their email or get to the Internet. Everything was fine yesterday, but since this morning they have had no connectivity. Of course, they haven't done anything to or changed their computer at all! No one else on the network appears to be affected.

If the computer is otherwise running normally, the first step should always be to run an `ipconfig` command to look at the IP address configured on the system. More often than not, the user will report back that their IP address is "169 dot 254 dot something dot something." The last two somethings don't really matter—it's the first two numbers that should have your attention. It's APIPA.

Knowing that the computer is a DHCP client, you know that it's not connecting to the DHCP server for some reason. After getting to the workstation, check the easy stuff first. Are the cables plugged in (if it's wired)? Are there lights on the NIC? Even if they appear to be plugged in, unplug and reconnect them. If that doesn't work, try a different cable. Those simple steps will solve the vast majority of these types of problems. If not, then it's on to more advanced troubleshooting steps! (More TCP/IP troubleshooting is covered in Volume 1 Chapter 12.)

IPv6

The present incarnation of TCP/IP that is used on the Internet was originally developed in 1973. Considering how fast technology evolves, it's pretty amazing to think that the protocol still enjoys immense popularity about 50 years later. This version is known as *IPv4*.

IPv6 is part of Exam Objective 2.6. CompTIA doesn't provide any details under IPv6, but be able to recognize some basic IPv6 addresses. Also know the differences between unicast, multicast, anycast, and link local addresses. There are a few problems with IPv4, though. One is that we're essentially out of available network addresses, and the other is that TCP/IP can be somewhat tricky to configure.

If you've dealt with configuring custom subnet masks, you may nod your head at the configuration part, but you might be wondering how we could run out of addresses. After all, IPv4 has 32 bits of addressing space, which allows for nearly 4.3 billion addresses! With the way it's structured, only about 250 million of those addresses are actually usable, and all of those are pretty much spoken for.

A new version of TCP/IP has been developed, called *IPv6*. Instead of a 32-bit address, it provides for 128-bit addresses. That provides for 3.4×10^{38} addresses, which theoretically should be more than enough that they will never run out globally. (Famous last words, right?)

IPv6 also has many standard features that are optional (but useful) in IPv4. While the addresses may be more difficult to remember, the automatic configuration and enhanced flexibility make the new version sparkle compared to the old one. Best of all, it's backward-compatible with and can run on the computer at the same time as IPv4, so networks can migrate to IPv6 without a complete restructure.

Understanding IPv6 Addressing

Understanding the IPv6 addressing scheme is probably the most challenging part of the protocol enhancement. The first thing you'll notice is that, of course, the address space is longer. The second is that IPv6 uses hexadecimal notation instead of the familiar dotted decimal of IPv4. Its 128-bit address structure looks something like what is shown in Figure 7.11.

FIGURE 7.11 IPv6 address

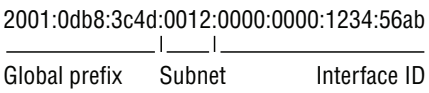

The new address is composed of eight 16-bit fields, each represented by four hexadecimal digits and separated by colons. The letters in an IPv6 address are not case sensitive. IPv6 uses three types of addresses: *unicast*, *anycast*, and *multicast*. A unicast address identifies a single

node on the network. An anycast address refers to one that has been assigned to multiple nodes. A packet addressed to an anycast address will be delivered to the closest node. Sometimes you will hear this referred to as one-to-nearest addressing. Finally, a multicast address is one used by multiple hosts, and is used to communicate to groups of computers. IPv6 does not employ broadcast addresses. Multicasts handle that functionality. Unlike IPv4, where each network interface can have only one IP address, in IPv6, each network interface can be assigned one or more addresses.

Just by looking at unicast and anycast addresses, it's impossible to tell the difference between them. Their structure is the same; it's their functionality that's different. The first four fields, or 64 bits, refer to the network and subnetwork. The last four fields are the interface ID, which is analogous to the host portion of the IPv4 address. Typically, the first 56 bits within the address are the routing (or global) prefix, and the next 8 bits refer to the subnet ID. It's also possible to have shorter routing prefixes, though, such as 48 bits, meaning that the subnet ID will be longer.

The Interface ID portion of the address can be created in one of four ways. It can be created automatically using the interface's MAC address, procured from a DHCPv6 server, assigned randomly, or configured manually.

Multicast addresses can take different forms. All multicast addresses use the first 8 bits as the prefix.

Working with IPv6 Addresses

In IPv4, the subnet mask determines the length of the network portion of the address. The network address was often written in an abbreviated form, such as 169.254.0.0/16. The /16 indicates that the first 16 bits are for the network portion and that corresponds to a subnet mask of 255.255.0.0. While IPv6 doesn't use a subnet mask, the same convention for stating the network length holds true. An IPv6 network address could be written as 2001:db8:3c4d::/48. The number after the slash indicates how many bits are in the routing prefix.

Because the addresses are quite long, there are a few ways that you can write them in shorthand; in the world of IPv6, it's all about eliminating extra zeroes. For example, take the address 2001:0db8:3c4d:0012:0000:0000:1234:56ab. The first common way to shorten it is to remove all of the leading zeroes. Thus it could also be written as 2001:db8:3c4d:12:0:0:1234:56ab. The second accepted shortcut is to replace consecutive groups of zeroes with a double colon. So now the example address becomes 2001:db8:3c4d:12::1234:56ab. It's still long, but not quite as long as the original address.

The double-colon shortcut can be used only once in an address. For example, in the 2001:db8:3c4d:12::1234:56ab address, you can count the number of fields (six) and know that the double colon represents two fields of all zeroes. If, for example, you tried to write an address like 2001::1ab4::5468, you would have a big problem. You would know that there are five fields of zeroes, but you would have no way to identify where exactly the 1ab4 portion of the address falls in relation to the all-zero fields.

A common occurrence today is a mixed IPv4-IPv6 network. As mentioned earlier, IPv6 is backward-compatible. In the address space, this is accomplished by setting the first 80 bits to 0, the next 16 bits to 1, and the final 32 bits to the IPv4 address. In IPv6 format, the IPv4 address looks something like ::ffff:c0a8:173. You will often see the same address written as ::ffff:192.168.1.115 to enable easy identification of the IPv4 address.

There are a few more addresses you need to be familiar with. In IPv4, the autoconfiguration (APIPA) address range was 169.254.0.0/16. IPv6 accomplishes the same task with the *link local* address fe80::/10. Every IPv6-enabled interface is required to have a link local address, and they are nonroutable. The IPv4 loopback address of 127.0.0.1 has been replaced with ::1/128 (typically written as just ::1). Global addresses (for Internet use) are 2000::/3, and multicast addresses are FF00::/8. Figure 7.12 shows the output of an `ipconfig` command, and you can see the IPv4 address configuration as well as the IPv6 link local address. Table 7.6 summarizes the IPv6 address ranges you should be familiar with.

TABLE 7.6 IPv6 address ranges

Address	Use
0:0:0:0:0:0:0:0	Equals ::, and is equivalent to 0.0.0.0 in IPv4. It usually means that the host is not configured.
0:0:0:0:0:0:0:1	Also written as ::1. Equivalent to the loopback (testing) address of 127.0.0.1 in IPv4.
2000::/3	Global unicast address range for use on the Internet.
FC00::/7	Unique local unicast address range.
FE80::/10	Link local unicast range.
FF00::/8	Multicast range.

FIGURE 7.12 `ipconfig` output with IPv4 and IPv6 addresses

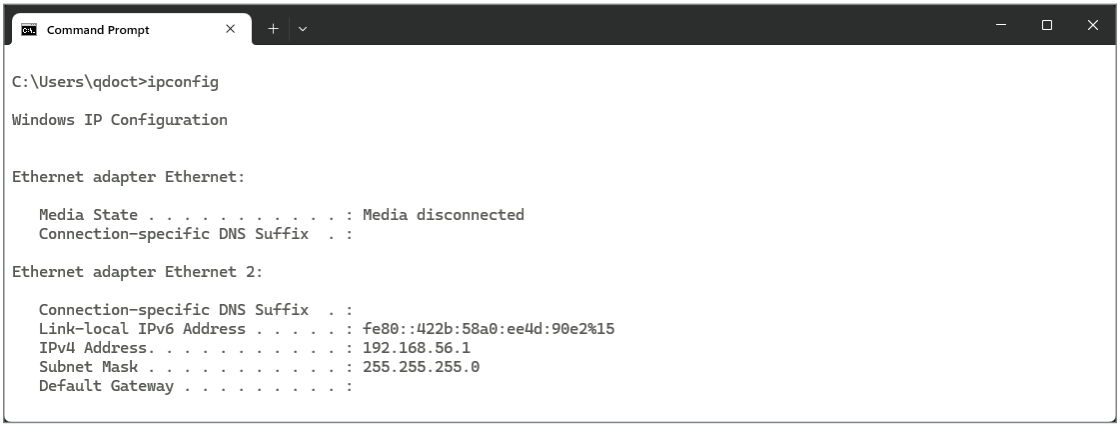
Command Prompt
C:\Users\qdoct>ipconfig
Windows IP Configuration
Ethernet adapter Ethernet:
Media State . . . . . : Media disconnected
Connection-specific DNS Suffix . :
Ethernet adapter Ethernet 2:
Connection-specific DNS Suffix . :
Link-local IPv6 Address . . . . . : fe80::422b:58a0:ee4d:90e2%15
IPv4 Address. . . . . : 192.168.56.1
Subnet Mask . . . . . : 255.255.255.0
Default Gateway . . . . . :

Understanding Virtual Networks

As you learned earlier in this chapter, the subnet mask on an IPv4 network determines the network address. Said differently, it's the mechanism by which networks are defined. Computers configured to be on different networks talk to each other through a router, which sends packets from one network to another. Therefore, the router is the physical device that divides logical networks from each other. In addition to physical and logical networks, one additional term you need to be familiar with is the virtual network. There are two types of virtual networks we'll cover here: virtual local area networks and virtual private networks.

Exam Objective 2.4, "Explain common networking configuration concepts," lists VLANs and VPNs as test objectives. Be sure to know the purposes for each and the differences between them for the exam. (VPNs will also be tested in exam 220-1202, and we talk about them again in Volume 1 Chapter 16.) The exam objective also lists DNS and DHCP, and we will cover those in depth in Volume 1 Chapter 8.

Virtual Local Area Networks

One of the limitations of typical routed network configurations is that computers on the same side of the router can't easily be broken into multiple networks and still communicate with each other. This is because if a sending computer knows that the destination IP address is on another network, it sends its data directly to the router—its default gateway. Other computers on the physical segment will ignore the message because it's not addressed to them. The router then takes a look at the real destination address and sends it out one of its ports, other than the one it came in on, to reach the destination network.

The *virtual local area network (VLAN)* is designed to help segment physical networks into multiple logical (virtual) networks. You may recall from Volume 1 Chapter 6 that VLANs are created by using a managed switch. The switch uses *Spanning Tree Protocol (STP)* to manage configurations and to ensure that there are no infinite network loops. (That's when data gets sent out and bounces between two or more switches, never getting to a destination. Loops are bad.) A VLAN can provide the following benefits:

Broadcast traffic is reduced. Physical network segments can be logically subdivided, reducing broadcast traffic and speeding network performance.

Security is increased. Computers on the same physical network can be isolated from each other to provide an additional layer of security. For example, imagine that the Research and Development team is working on several secret projects. R&D computers can be placed on one virtual segment, and the rest of the company on others. Traffic to and from other R&D computers won't be intermingled with other office traffic.

Computers in multiple locations can belong to the same VLAN. This is one major thing that routers can't do with subnetting. With multiple switches configured

appropriately, computers at different physical locations can be configured to be on the same VLAN.

Reconfiguring networks is easier. With VLANs, if someone moves physical desk locations, their VLAN membership can carry with them, so there is less network reconfiguration needed. Similarly, if someone doesn't move desk locations but changes jobs, they can be assigned to a new VLAN without needing to physically reconfigure the network.

Figure 7.13 shows two potential VLAN configurations. In the first one, computers on one switch are assigned to different VLANs. In the second, the concept is extended to include multiple switches.

FIGURE 7.13 Two VLAN configurations

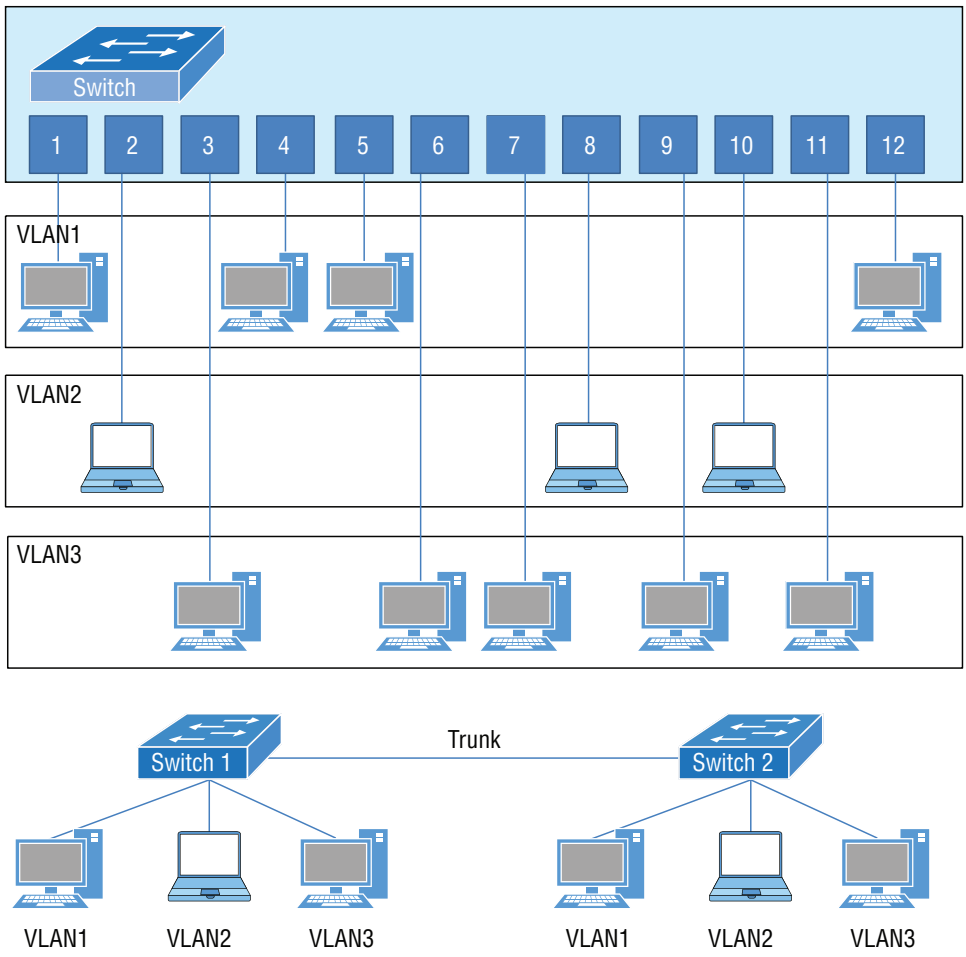

One of the questions often asked is, "What's the difference between a VLAN and a subnet?" First, let's look at the key similarity—both are capable of breaking up broadcast

domains on a network, which helps reduce network traffic. Also, if you are using both, the recommended configuration is that subnets and VLANs have a 1:1 relationship, one subnet per VLAN. You can configure multiple subnets to be on one VLAN—it's called a *super scope*—but it gets trickier to manage.

Beyond separating broadcast domains, VLANs and subnets are almost entirely different. Recall that VLANs are implemented on switches, and routers are needed to subnet. Consequently, VLANs work at Layer 2 of the OSI model and deal with physical MAC addresses. Routers work at Layer 3, and they work with logical IP addresses.

Virtual Private Networks

As networks grow beyond simple physical limitations (such as an office or a building) to include clients from all over the world, the need to secure data across public connections becomes paramount. You've already learned about HTTPS, which encrypts data in transit between a web server and a web client. That's a very specific application—it's only used when two remote computers are communicating using web client and web server software. Other uses, such as a different type of app or the need to secure all data in transit being sent across public networks such as the Internet, require different solutions. Enter the *virtual private network (VPN)*.

A VPN is a secure (private) network connection that occurs through a public network. The private network connection provides data encryption over an otherwise unsecure environment. VPNs can be used to connect LANs together across the Internet or other public networks, or they can be used to connect individual users to a corporate network. This is a great option for users who work from home or travel for work. With a VPN, the remote end appears to be connected to the network as if it were connected locally. Conceptually, this looks like the example shown in Figure 7.14.

FIGURE 7.14 A VPN

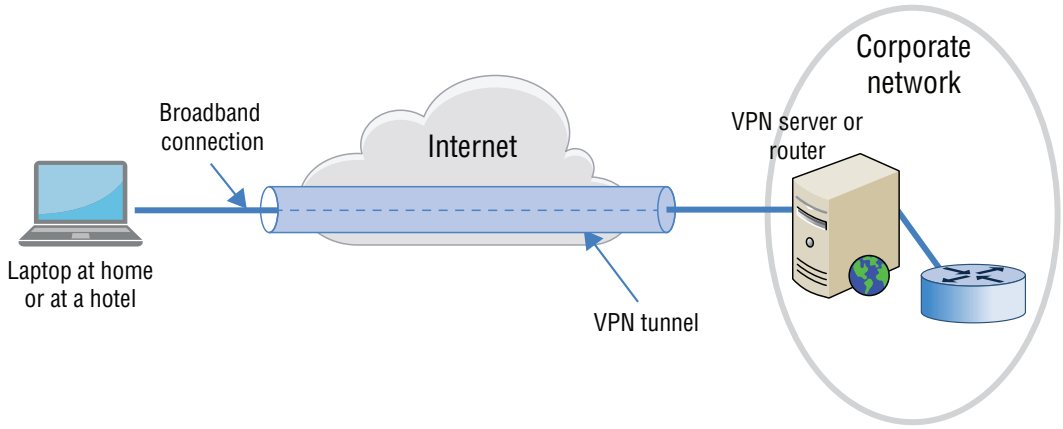

How VPNs Work

VPNs use the concept of network tunneling to secure the connection and encrypt data between two remote networks. Tunneling sends private data across a public network by placing (encapsulating) that data into other packets. Before we dive into what that means, we should clarify that there are a few different types of VPNs available. They differ based on which protocols they use to secure the connection, and consequently which level of the OSI model they operate at.

SSL/TLS VPNs In the "Port 443—Hypertext Transfer Protocol Secure (HTTPS)" section earlier in this chapter, you were introduced to Secure Sockets Layer and Transport Layer Security. Their jobs are to encrypt the HTTPS traffic between the web server and client. SSL and TLS can encrypt other types of data as well, not just web traffic. If a VPN is running SSL/TLS, the data is being encrypted at the Application layer—Layer 7 of the OSI model. Therefore, an SSL/TLS VPN is typically responsible for one specific application only. It does not encrypt the entire networked data stream between two remote computers.

IPSec VPNs The other major type of VPN is one secured by the *Internet Protocol Security (IPSec)* protocol. IPSec works at OSI Layer 3 and can encrypt all packets being sent between two computers across a public network. There are two protocols that are part of IPSec that help with the process. The first is called Authentication Header (AH), and it ensures data integrity and the origin of the data. The second is Encapsulating Security Payload (ESP), and that's the protocol specifically in charge of encryption.

From here on, we'll focus on IPSec VPNs to explain how VPNs work.

The bundle of data at OSI Layer 3 is referred to as a packet. Packets consist of three different parts:

- An IP header, which among other things contains the sending and receiving IP addresses and which protocol is being used (TCP or UDP)
- The data portion, or payload, which is the actual information being sent between computers
- An optional trailer, which might include a signal to indicate the end of a packet or error detection information

In normal, unencrypted transmission, the data is sent to the transmission medium, such as the network cable, and on to the receiving device. If another computer with the right hardware and software, such as a packet sniffer, intercepts the message, the data is all there in plain text and easily read (provided, of course, that no other encryption was used). Yes, a packet is small and only a fraction of an entire message, but the sniffer can use data in the IP header to put it all together.

When using a VPN, ESP interjects and adds its own header information and also includes a trailer to encrypt the transmission. Only a device with the proper decoding key can unencrypt the data at the receiving end.

There are two types of IPSec VPNs: transport mode and tunnel mode. In transport mode, all data is encrypted but the original IP header is left intact. Most client-server VPNs, such as those used by people working from home or traveling, use transport mode. In tunnel mode, the original IP header is encrypted as well. This is more often used when the entire network is hiding behind a translator such as NAT. All of this is illustrated in Figure 7.15.

FIGURE 7.15 IPSec VPNs

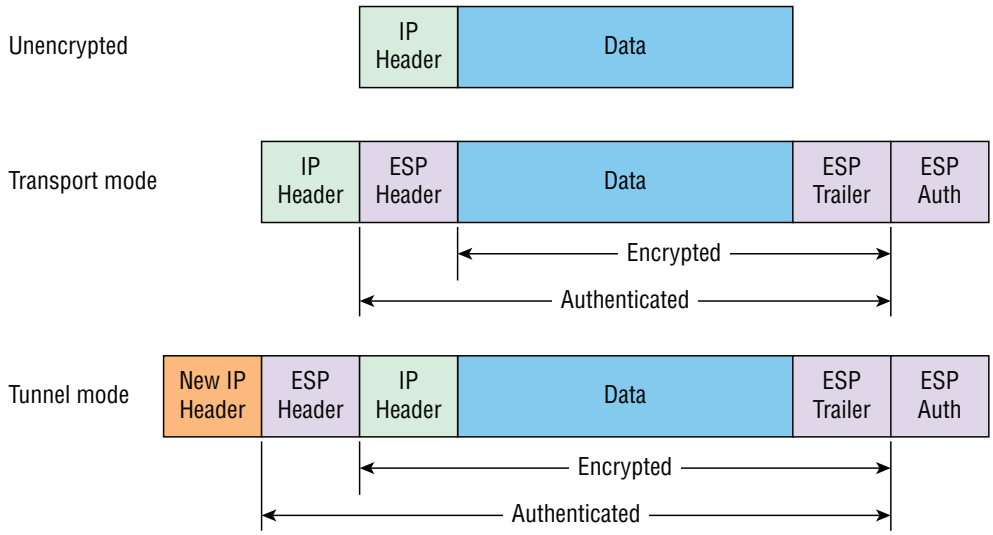

Configuring VPN clients

Users wanting to employ a VPN will connect to the Internet using their normal connection methods, such as Wi-Fi. Then they will use a software program to initiate the VPN connection. There are a multitude of VPN clients available for purchase on the Internet, and Windows has its own built-in client as well. Exercise 7.2 shows you how to configure it.

EXERCISE 7.2

Setting Up Windows 11 VPN

This exercise has you configure the Windows 11 VPN client.

1. Open Settings by clicking Start, then Settings.
2. From the left-side navigation menu, choose Network & Internet.
3. Click the VPN box.
4. Next to VPN Connections, click Add VPN. You will get a screen like the one shown in Figure 7.16.

EXERCISE 7.2 *(continued)*

FIGURE 7.16 Adding a VPN

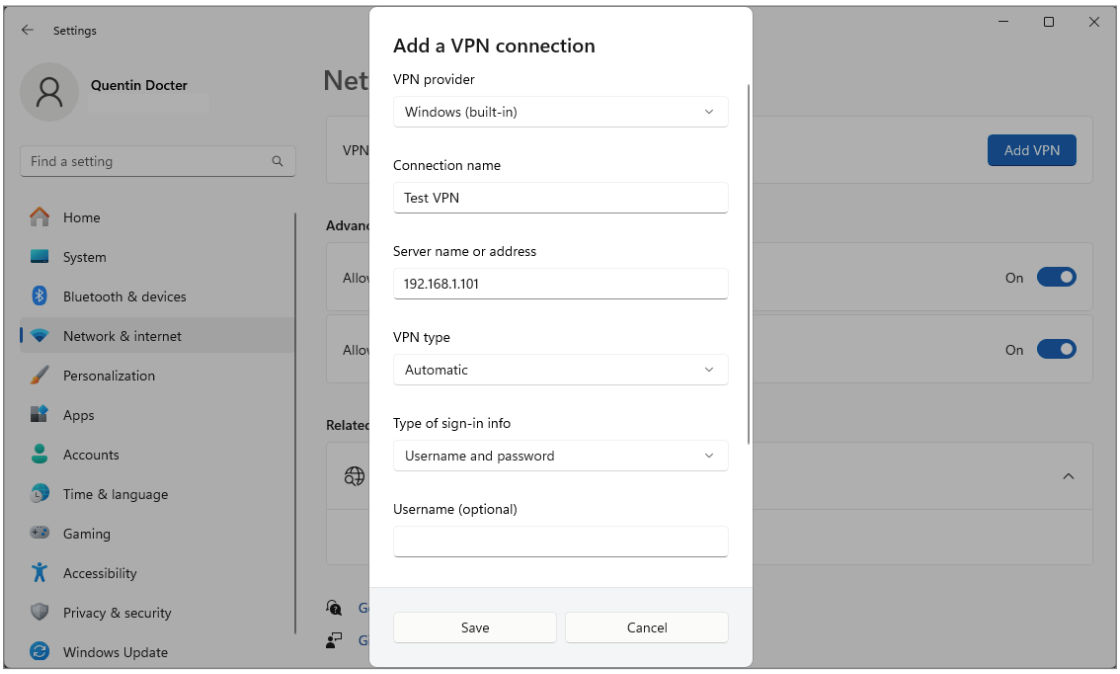

5. If you have real VPN server configuration information to add, use it. Otherwise, you can make up a name and IP address as shown in Figure 7.16.

   The VPN type setting lets you configure specific security protocols. The type of sign-in info has choices such as username and password and a one-time security token. The correct settings depend on how the VPN server is configured.

6. Click Save. The VPN will now appear as shown in Figure 7.17.

FIGURE 7.17 Test VPN added

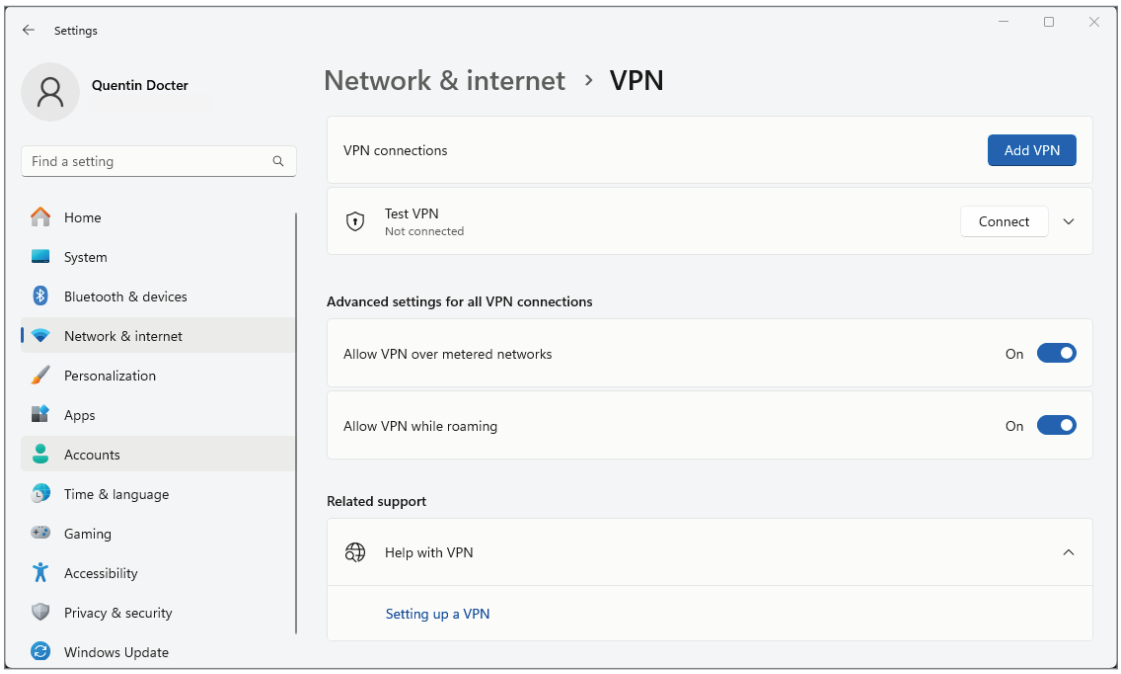

7. To connect to the VPN, click the Connect button.
8. You will receive a pop-up window asking for your username and password. If you put in real VPN server data, you can test the connection now. Otherwise, just cancel.
9. (optional) To remove the VPN, click the down arrow next to the Connect button and choose Remove.

Summary

In this chapter, you learned about the protocol suite used on the Internet, TCP/IP. It's by far the most common protocol in worldwide use today. We started with TCP/IP structure. It's a modular suite that follows the DoD model, with different protocols performing unique tasks at each layer. We looked at individual protocols and their functions at the Internet, Host-to-host, and Process/Application layers. We also discussed ports and well-known port numbers for common protocols.

Next you learned about IP addressing. We started with a brief tutorial on converting binary numbers to decimal to make them easier to read. Then we looked at the different address classes, CIDR, public versus private IP addresses, NAT, and APIPA. Each of these concepts play a unique role in managing TCP/IP on your network.

Next, you learned about the current generation of TCP/IP, IPv6. We talked about the seemingly infinite number of addresses as well as the fact that addresses are written in hexadecimal, which might take some getting used to—even for experienced technicians. We then looked at working with IPv6 addresses, including shorthand notation and special addresses to be aware of.

Finally, we finished the chapter by looking at two types of virtual networks: VLANs and VPNs. VLANs create logical network groups using a switch. VPNs create a private, encrypted connection across a public network.

Exam Essentials

Understand how IPv4 addressing works. IP addresses are 32-bit addresses written as four octets in dotted-decimal notation, such as 192.168.5.18. To communicate on an IP network, a host also needs a subnet mask, which may look something like 255.255.255.0. If the host needs to communicate outside the local network, it also needs a default gateway, which is normally the internal address of the router.

Addresses can be static (manual) or dynamic (from a DHCP server). If a DHCP server is not available, a network client may use an APIPA address with the format 169.254.x.x. Be able to identify IP address classes. Know how to identify Class A, B, and C IP addresses. Class A addresses will have a first octet in the 1 to 126 range. B is from 128 to 191, and C is from 192 to 223.

Understand the differences between TCP and UDP. TCP is a connection-based protocol that attempts to guarantee delivery. UDP is connectionless, which makes it a bit faster, but it doesn't guarantee packet delivery.

Know common TCP/IP ports. Some common protocol and port pairings that you should know are FTP (20 and 21), SSH (22), Telnet (23), SMTP (25), DNS (53), DHCP (67, 68), HTTP (80), POP3 (110), NetBIOS/NetBT (137-139), IMAP (143), LDAP (389), HTTPS (443), SMB/CIFS (445), and RDP (3389).

Know the difference between public and private IP addresses. Public addresses are routable on the Internet and must be globally unique. Private addresses are not routable on the Internet and are used behind a NAT device.

Know the private IP address ranges. Private IP addresses will be in one of three ranges: 10.0.0.0/8, 172.16.0.0/12, or 192.168.0.0/16.

Know what NAT does. Network Address Translation (NAT) translates private, non-routable IP addresses into public IP addresses. It allows computers on a private network to access the Internet.

Know about the APIPA range. IP addresses in the 169.254.0.0/16 range are APIPA addresses.

Understand how IPv6 addressing works. IPv6 addresses are 128-bit addresses written as eight fields of four hexadecimal characters, such as 2001:0db8:3c4d:0012: 0000: 0000:1234:56ab. Using shorthand conventions, this address can also be written as 2001:db8:3c4d:12::1234:56ab.

Addresses can be static or dynamic. APIPA does not exist in IPv6 but has been replaced by a link local address.

Know the difference between unicast, anycast, and multicast in IPv6. Unicast addresses are for a single node on the network. Anycast can represent a small group of systems. An anycast message will be delivered to the closest node. Multicast messages are delivered to all computers within a group.

Recognize the special classes of IPv6 addresses. The loopback address is ::1. Global unicast addresses are in the 2000::/3 range. Unique local unicast addresses are in the FC00::/7 range, link local addresses are FE80::/10, and FF00::/8 addresses are multicast.

Understand the differences between a VLAN and a VPN. A virtual local area network (VLAN) is a logical network configured through a managed switch. A virtual private network (VPN) is a secure point-to-point connection over a public network.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. You are setting up clients to use the new corporate VPN. The VPN will encrypt all data in transit between employees who work from home and the corporate office. Which protocol will secure the VPN?
   - **A.** HTTPS
   - **B.** IPSec
   - **C.** TLS
   - **D.** SSL
2. Which of the following protocols was designed to access information stored in an information directory, such as employee phonebooks?
   - **A.** FTP
   - **B.** LDAP
   - **C.** RDP
   - **D.** CIFS
3. Which TCP/IP protocol uses port 80?
   - **A.** HTTP
   - **B.** HTTPS
   - **C.** Telnet
   - **D.** POP3
4. What is the maximum number of IPv6 addresses that can be assigned to one IPv6 interface?
   - **A.** One (unicast)
   - **B.** Two (unicast and anycast)
   - **C.** Three (unicast, anycast, and multicast)
   - **D.** None of the above
5. Which of the following are valid examples of IPv6 addresses? (Choose all that apply.)
   - **A.** 2001:0db8:3c4d:0012:0000:0000:1234:56ab
   - **B.** ::ffff:c0a8:173
   - **C.** 2001:db8:3c4d:12::1234:56ab
   - **D.** 2001::1ab4::5468
6. Which of the following IP addresses would not be valid for a DNS server on the Internet?
   - **A.** 10.25.11.33
   - **B.** 18.33.66.254
   - **C.** 155.118.63.11
   - **D.** 192.186.12.2

7. The workstations on your network are configured to use a DHCP server. One of the workstations can't communicate with other computers. Its IP address is 169.254.1.18. What could be the problem?
   - **A.** The subnet mask is wrong.
   - **B.** It has a private IP address.
   - **C.** The default gateway is wrong.
   - **D.** It can't reach the DHCP server.

8. Which of the following protocols is responsible for sending email?
   - **A.** IMAP4
   - **B.** POP3
   - **C.** SMTP
   - **D.** SNMP

9. What port does the RDP protocol work on?
   - **A.** 53
   - **B.** 143
   - **C.** 389
   - **D.** 3389

10. Which of the following solutions can help reduce broadcast traffic received by TCP/IP-based computers on a network?
    - **A.** DNS
    - **B.** DHCP
    - **C.** VPN
    - **D.** VLAN

11. What are advantages that TCP has over UDP? (Choose two.)
    - **A.** Acknowledged delivery
    - **B.** Faster delivery
    - **C.** Lower overhead
    - **D.** Virtual circuits

12. Your friend is concerned about the security of making an online purchase. What should you tell her to look for in the address bar of the web browser?
    - **A.** HTTP
    - **B.** HTTPS
    - **C.** SSH
    - **D.** TLS

13. You are manually configuring a TCP/IP host. Another administrator gives you the router's IP address. What is the TCP/IP term for this?

A. Default gateway

B. Subnet mask

C. DNS server

D. DHCP server

14. Your network is running IPv4. Which of the configuration options are mandatory for your host to communicate on the network? (Choose all that apply.)

A. IP address

B. Subnet mask

C. Default gateway

D. DNS server address

15. Which of the following protocols is used for secure retrieval of email?

A. SMTP

B. LDAP

C. POP3

D. IMAP4

16. Which of the following is the name of a secure point-to-point connection made over a public network?

A. NAT

B. APIPA

C. VLAN

D. VPN

17. Which network protocol is used for file and printer sharing and which port does it use?

A. CIFS, 445

B. RDP, 3389

C. CIFS, 3389

D. RDP, 445

18. Which of the following is an IPv6 broadcast address?

A. ::1

B. FE80::

C. FF00::

D. ::FFFF

E. None of the above

19. Which of the following protocols can help an application store and manage user information on the network?

A. LDAP

B. SMTP

C. NetBT

D. SMB

20. Which of the following protocols is responsible for resolving IP addresses to hardware addresses?

A. DNS

B. DHCP

C. ARP

D. RARP

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answers compare to the authors', refer to Appendix B.

Match the following protocols (services) below to their respective ports in the table:

DHCP, DNS, FTP, HTTP, HTTPS, IMAP, LDAP, NetBIOS/NetBT, POP3, RDP, SMB/CIFS, SMTP, SSH, Telnet

Protocol (service)	Port(s)
	20–21
	22
	23
	25
	53
	67, 68
	80
	110
	137-139
	143
	389
	443
	445
	3389

Chapter 8

Network Services

THE FOLLOWING COMPTIA A+ 220-1201 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

- ✓ **2.0 Networking**
- ✓ **2.3 Summarize services provided by networked hosts.**
- ✓ **Server roles**
  - DNS
  - DHCP
  - Fileshare
  - Print servers
  - Mail servers
  - Syslog
  - Web servers
  - Authentication, Authorization, and Accounting (AAA)
  - Database servers
  - Network Time Protocol (NTP)
- ✓ **Internet appliances**
  - Spam gateways
  - Unified threat management (UTM)
  - Load balancers
  - Proxy servers
- ✓ **Legacy/embedded systems**
  - Supervisory control and data acquisition (SCADA)
- ✓ **Internet of Things (IoT) devices**
- ✓ **2.4 Explain common network configuration concepts.**
  - DNS

        - A
        - AAAA
        - Canonical name (CNAME)
        - Mail exchanger (MX)
        - Text (TXT)
            - Spam management
                - DomainKeys Identified Mail (DKIM)
                - Sender Policy Framework (SPF)
                - Domain-based Message Authentication, Reporting, and Conformance (DMARC)
    - DHCP
        - Leases
        - Reservations
        - Scope
        - Exclusions

Networks are often complicated structures. When users get on a network, they have expectations that features will be available to them. Most of the time they are unaware of the underlying infrastructure. As long as what they want gets delivered, they are content. In client-server networks, which you learned about in Volume 1 Chapter 6, "Networking Fundamentals," there are one or more servers that play unique roles in fulfilling client requests.

From the server side, the features users need and depend on are often implemented as services. This can include things like providing easy access to Internet servers by using their website names, storing files, hosting printers, managing email, and handling security. In this chapter, we will talk about some key network services you need to be familiar with as a technician. Servers provide some services, and stand-alone security devices or Internet appliances provide others. Often, it's a combination of devices that complete the picture that clients on the network need to see. This chapter explores several critical services in depth.

Understanding Network Services

As you learned in Volume 1 Chapter 6, networks are made up of multiple types of devices operating together. There are different types of networks, such as peer-to-peer, client-server, and cloud-based, and they are categorized based on the types of devices they support and where they are located. Very simple, small networks can be peer-to-peer, where there are no dedicated servers. However, most networks that you encounter in the business world will have at least one server, and enterprise networks can easily have hundreds or even thousands of servers.

New computer technicians won't be expected to manage large server farms by themselves, but they should be aware of the types of servers and other devices that will be on the network and their basic functions. Experienced technicians may be in charge of one or more servers, and they will need to be intimately familiar with their inner workings. As you gain more experience, you will find that there are advanced certifications in the market to prove your knowledge of servers and show off your skills to potential employers. For the A+ exam, you will need to know various server roles, the features of a few Internet security appliances, the impact of legacy and embedded systems, and the features of Internet of Things (IoT) devices.

Server Roles

Servers come in many shapes and sizes, and they can perform many different roles on a network. Servers are generally named for the type of service they provide, such as a web server or a print server. They help improve network security and ease administration by

centralizing control of resources and security; without servers, every user would need to manage their own security and resource sharing. Not everyone has the technical ability to do that, and even if they did, those types of responsibilities might not be part of what they are being asked to deliver at work. Servers can also provide features such as load balancing and increased reliability. Figure 8.1 shows a screenshot from Windows Server 2025's Server Manager dashboard. It gives you an example of different types of roles this server could play. Don't worry about memorizing the list, or knowing how to use Server Manager—the figure is intended to just give you a flavor for what the server environment looks like.

FIGURE 8.1 Server roles

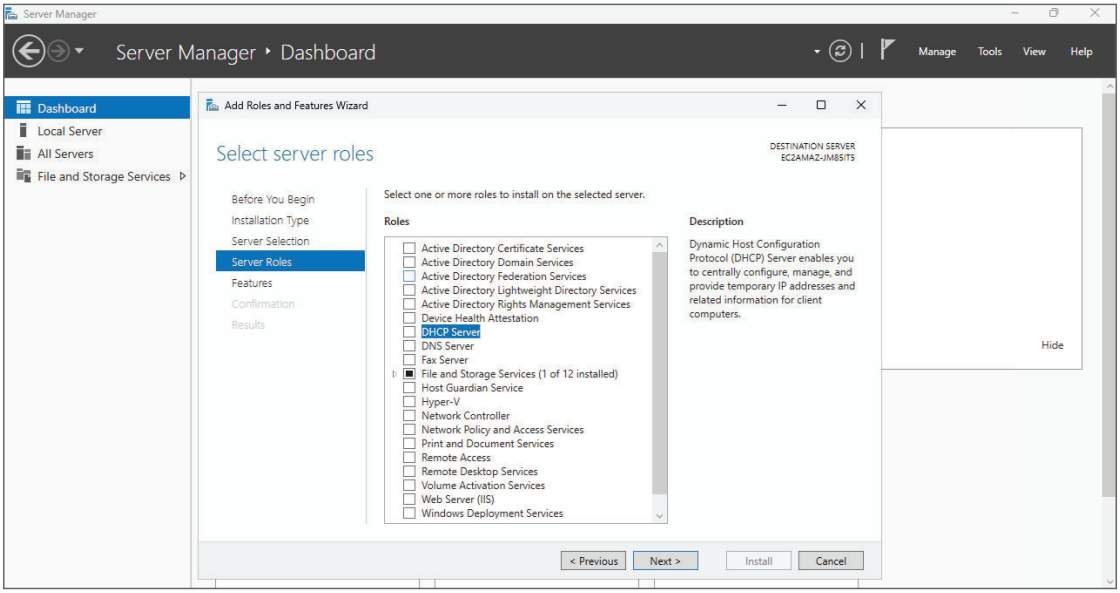

Most server operating systems, such as the Windows Server series, come with built-in software that allows them to play one or more server roles. The administrator simply needs to enable the role (such as DHCP or DNS) and configure it properly.

Some servers are dedicated to a specific task, such as hosting websites, and they are called *dedicated servers*. *Nondedicated servers* may perform multiple tasks, such as hosting a website and serving as the administrator's daily workstation. Situations like this are often not ideal because the system needs more resources to support everything it needs to do. Imagine that you are the user of that computer and there is heavy website traffic. Your system could slow down to the point where it's difficult to get anything done. In addition, that kind of setup could introduce additional security risks. Servers can also perform multiple server-specific roles at the same time, such as hosting websites and providing file and print services.

As you read through the descriptions of server roles, you will see that it makes more sense to combine some services than it does to combine others.

One important decision network architects need to make when thinking about designing a network is where to place the server or servers. In Volume 1 Chapter 6, we introduced the concept of a *screened subnet* (formerly called a *demilitarized zone [DMZ]*), which is a network separated from the internal network by a firewall but also protected from the Internet by a firewall. Figure 8.2 shows two examples of screened subnets.

FIGURE 8.2 Screened subnets

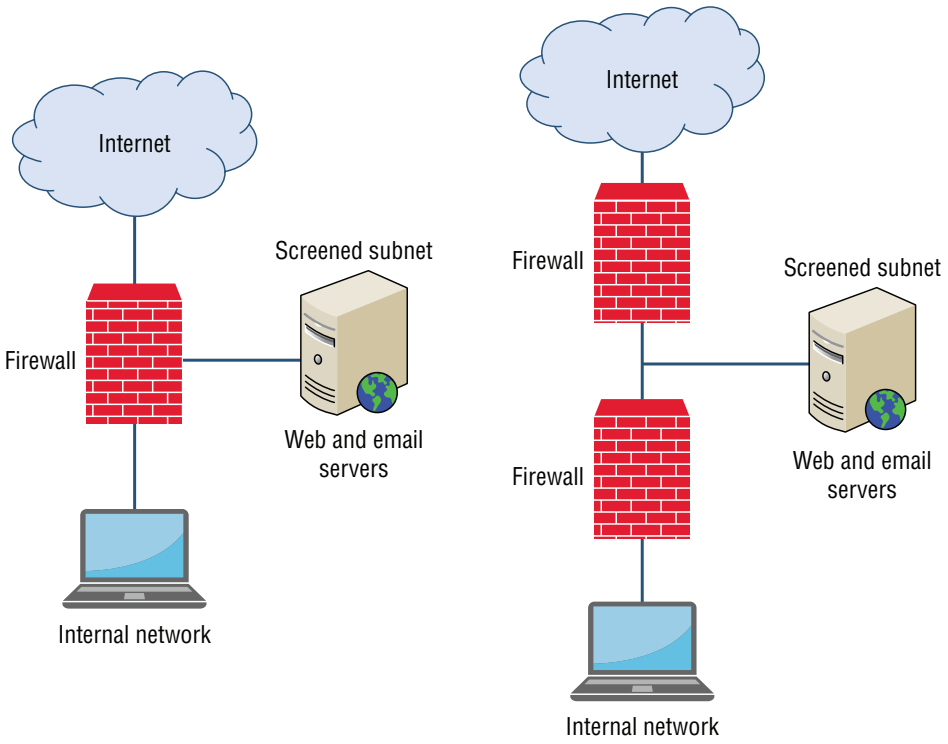

The network on the left in Figure 8.2 shows a screened subnet managed by what is called a *three-pronged firewall*—so named because it has three network interfaces. You can see that the firewall has an Internet interface, an internal interface, and a screened subnet interface. Another common configuration is a two-pronged firewall (shown on the right of Figure 8.2), with the screened subnet in the middle of the two. In that setup, if an attacker wants to get to your internal network, they have to go through two separate firewalls. Hopefully, you have your firewalls set up to make that task difficult for potential hackers!

In Figure 8.2, you see that the web and mail servers are in the screened subnet and not on the internal network. This configuration can make it easier to manage the network but still provide great security. As a rule of thumb, any server that needs to be accessed by the outside world should be in the screened subnet, and any server that does not need to be accessed from the Internet should be on the internal network, which is more secure. By the way, servers can play the role of firewalls, too. It's not, however, on the list of objectives as a server role, and in practice it's best to separate other server roles from firewalls. In other words, if you intend to use a server as a firewall, then don't use it for any other types of services. Having services on the firewall itself just makes it easier for hackers to get to. There's no sense in making things easier for them. Now it's time to talk about specific server roles on a network.

DNS

We briefly introduced *Domain Name System (DNS)* servers in Volume 1 Chapter 7, "Introduction to TCP/IP." There you learned that they resolve host names to IP addresses, such as resolving `www.google.com` to 2607:f8b0:4002:c0c::69 so that communication can begin, and that the DNS protocol uses UDP or TCP on port 53. Now we will dive into the inner workings of a DNS server.

DNS is a very important topic and is listed twice in the A+ 220-1201 objectives. Exam Objective 2.3, "Summarize services provided by networked hosts," lists DNS as a server role to know. Exam Objective 2.4, "Explain common network configuration concepts," gets into a lot more detail on what you need to know. We will add additional tips throughout this section to highlight those specifics.

That DNS has only one function on the network, to resolve hostnames to IP addresses, makes it sound pretty simple. Really though it's more complicated than it appears on first blush, and it has profound implications.

Most people use the Internet daily, and visit multiple websites. You open your browser, and in the address bar, you type the name of your favorite website, something like `www.comptia.com`, and press Enter. The first question your computer asks is, "Who is that?" Your machine requires an IP address to connect to the website. The DNS server provides the answer, "That is 2606:4700:3033::6815:4b4c." Now that your computer knows the address of the website you want, it's able to traverse the Internet to find it.

Each DNS server has a database where it stores hostname-to-IP-address pairs. If the DNS server does not know the address of the host you are seeking, it has the ability to query other DNS servers to help answer the request.

Think about the implications of that for just a minute. We all probably Google something several times a day, but in all honesty how many of us know its IP address? It's certainly not something we are likely to have memorized. Much less, how could you possibly memorize

the IP addresses of all the websites that you regularly visit? Because of DNS, it's easy to find resources. Whether you want to find Coca-Cola, Toyota, Amazon, or thousands of other companies, it's usually pretty easy to figure out how. Type in the name with a .com on the end of it, and you're usually right. The only reason this is successful is because DNS is there to perform resolution of that name to the corresponding IP address.

DNS works the same way on an intranet (a local network not attached to the Internet) as it does on the Internet. The only difference is that instead of helping you find `www.google.com`, it may help you find Jen's print server or Joe's file server. From a client-side perspective, all you need to do is configure the host with the address of a legitimate DNS server and you should be good to go.

The DNS Server

If a company wants to host its own website, it also needs to maintain two public *DNS servers* with information on how to get to the website. (Two servers are required for redundancy.) An advantage of using ISPs or web hosting companies to host the website is that they are then also responsible for managing the DNS servers.

Each DNS server has a database, called a *zone file*, which maintains records of hostname to IP address mappings. Within a zone file, you will see information that looks something like this:

```
mydomain.com. IN   SOA   ns.mydomain.com.    ;Start of Authority record
mydomain.com. IN   NS    ns.mydomain.com.    ;name server for mydomain.com
mydomain.com. IN   MX    mail.mydomain.com. ;mail server for mydomain.com
mydomain.com. IN   A     192.168.1.25        ;IPv4 address for mydomain.com
              IN   AAAA  2001:db8:19::44     ;IPv6 address for mydomain.com
ns            IN   NS    192.168.1.2         ;IPv4 address for ns.mydomain.com
www           IN   CNAME mydomain.com.       ;www.mydomain.com is an
                                                 alias for mydomain.com
www2          IN   CNAME www                 ;www2.mydomain.com is another
                                                 alias for mydomain.com
mail          IN   A     192.168.1.26        ;IPv4 address for
                                                 mail.mydomain.com
```

Five columns of information are presented. From left to right, they are as follows:

- The name of the server or computer, for example, `www`.
- IN, which means Internet. (There are other options for this field, but for our purposes, we will focus on Internet.)
- The record type. This example has SOA, NS, MX, A, AAAA, and CNAME. Table 8.1 explains some of the common record types.
- The address of the computer. This can be IPv4 or IPv6.
- Comments, preceded by a semicolon. In a file like this, the computer disregards everything after a semicolon. It's used to make notes for the administrator without affecting functionality.

TABLE 8.1 Common DNS record types

Type	Meaning
SOA	Start of Authority. It signifies the authoritative DNS server for that zone.
NS	Name Server. It's the name or address of the DNS server for that zone.
MX	Mail Exchange. It's the name or address of the email server.
A	IPv4 host record.
AAAA	Called "quad A," it's the host record for IPv6 hosts.
CNAME	Canonical Name. It's an alias; it allows multiple names to be assigned to the same host or address.
TXT	Text record. Used to enter human-readable or machine-readable data. Today, text records are used primarily for email spam prevention and domain ownership verification.

Remember that the DNS record types A, AAAA, CNAME, MX, and TXT are specifically called out in Exam Objective 2.4.

The DNS server uses the zone file whenever a computer makes a query. For example, if you were to ask this DNS server, "Who is `mydomain.com`?" the response would be 192.168.1.25. If you ask it, "Who is `www.mydomain.com`?" it would look and see that `www` is an alias for `mydomain.com` and provide the same IP address.

If you are the DNS administrator for a network, you will be required to manage the zone file, including entering hostnames and IP addresses, as appropriate.

Spam Management

Email spam is a problem. The only people who don't agree with this are the spammers themselves or the folks who pay them to spam you! One of the tricks that spammers use is to spoof (or fake) the domain name they are sending emails from. DNS, through the use of TXT records, can help email servers determine if incoming messages are from a trusted source rather than a spoofed one.

The three TXT records in this section are sub-objectives of exam objective 2.4. Know that the three types of records exist, that they help protect against spam email, and the basics of how each one does its job. It is beyond the scope of the A+ exam to know how to configure each of them because it's a relatively advanced DNS management concept.

Three standards used to battle email spam are *Sender Policy Framework (SPF)*, *Domain Keys Identified Mail (DKIM)*, and *Domain-based Message Authentication, Reporting, and Conformance (DMARC)*. Each one can be used in a DNS TXT record to help thwart malicious users from using a company's domain name to send unauthorized emails.

SPF is the simplest of the three. It authenticates an email server based on its IP address. In an SPF TXT record, the administrator specifies all servers that are legitimate email senders for that domain, based on their IP addresses. Note that we're not referring to client computers sending email, but the addresses of the email servers that are legitimate. When a receiving email server gets a message, it sends a query back to the supposed sending server via the sending domain's return-path value, which is found in the email's headers. The sending domain's DNS servers will provide a list of approved senders' IP addresses. If the original sending machine's IP address is on the list, the email is accepted. If not, the email is rejected.

DKIM is a bit more involved, as it authenticates using encryption through a public-private key pair. Each email sent by the server includes a digital signature in the headers, which has been encrypted by the server's private key. When the receiving email server gets the message, it finds the server's registered public key, which is used to decrypt the message. If the key pair is incorrect, the message is flagged as a fake.

Encryption can quickly become a complicated topic. The A+ test won't get too deep into it—what we shared here about how the public-private key pair works should be good enough—but other exams such as Security+ definitely do. If you're interested in learning more about private and public keys, a good reference is `www.ssl.com/article/private-and-public-keys`.

DMARC builds on both SPF and DKIM and essentially combines them together into one framework. It's not an authentication method per se—rather, it allows a domain owner to decide how they want email from their domain to be handled if it fails either an SPF or a DKIM authentication. Options include doing nothing (letting the email through), quarantining the email (i.e., sending it to a spam folder), or rejecting the email. In addition, it allows the domain owner to see where emails that claim to come from their domain actually originate from.

DNS on the Internet

The Internet is really big—so big that there's no way one DNS server could possibly manage all of the computer name mappings out there. The creators of DNS anticipated this, and they designed it in a way that reduces potential issues. For example, let's say that you are looking for the website `www.wiley.com`. When the DNS server that your computer is configured to use is queried for a resolution, it will first check its zone file to see if it knows the IP address. If not, it then checks its cache to see if the record is in there. The cache is a temporary database of recently resolved names and IP addresses. If it still doesn't know the answer, it can query another DNS server asking for help. The first server it will ask is called a *root server*.

If you look back at the sample zone file shown earlier, you might notice that the first few rows contained `mydomain.com.` (the dot at the end, the "trailing dot," is intentional). The Internet name space is designed as a hierarchical structure, and the dot at the end is the broadest categorization, known as "the root." At the next level of the hierarchy are the top-level domains, such as `.com`, `.net`, `.edu`, `.jp`, and others. Below that are the second-level domains, like Google, Microsoft, and Yahoo. Below that there are subdomains (which are optional) and hostnames. Moving down, the levels in the hierarchy get more and more specific, until the name represents an exact host. Figure 8.3 shows an example.

FIGURE 8.3 Internet name hierarchy

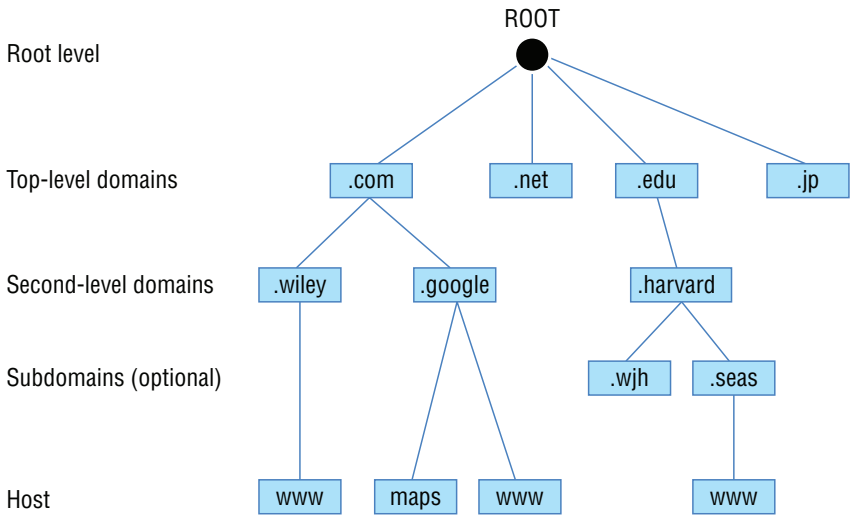

The Dot at the End of the Internet

You've probably used the Internet for quite some time, and you may be thinking, "Where did that trailing dot come from? I've never typed it into my browser." You're right.

As we said, the dot at the end represents the root. Without it, a domain name is not considered a *fully qualified domain name (FQDN)*, fit for Internet use. However, the dot is a convention that's implied when Internet browsers are used; users don't need to type it in. Even though it's omitted, the browser understands that it's technically looking for `www.yahoo.com.` and not `www.yahoo.com` with no terminal period. The Internet is an amazing place, isn't it?

There are 13 global root servers. (Technically speaking, there are 13 IP addresses for global root DNS servers. Each address represents several redundant servers located around the world.) All DNS servers should be configured to ask a root server for help. The root server will return the name of a top-level domain DNS server. The querying DNS server will

then ask that server for help. The process continues, as shown in Figure 8.4, until the querying DNS server finds a server that is able to resolve the name `www.wiley.com`. Then the querying DNS server will cache the resolved name so that subsequent lookups are faster. The length of time that the name is held in cache is configurable by the DNS administrator.

FIGURE 8.4 The DNS name resolution process

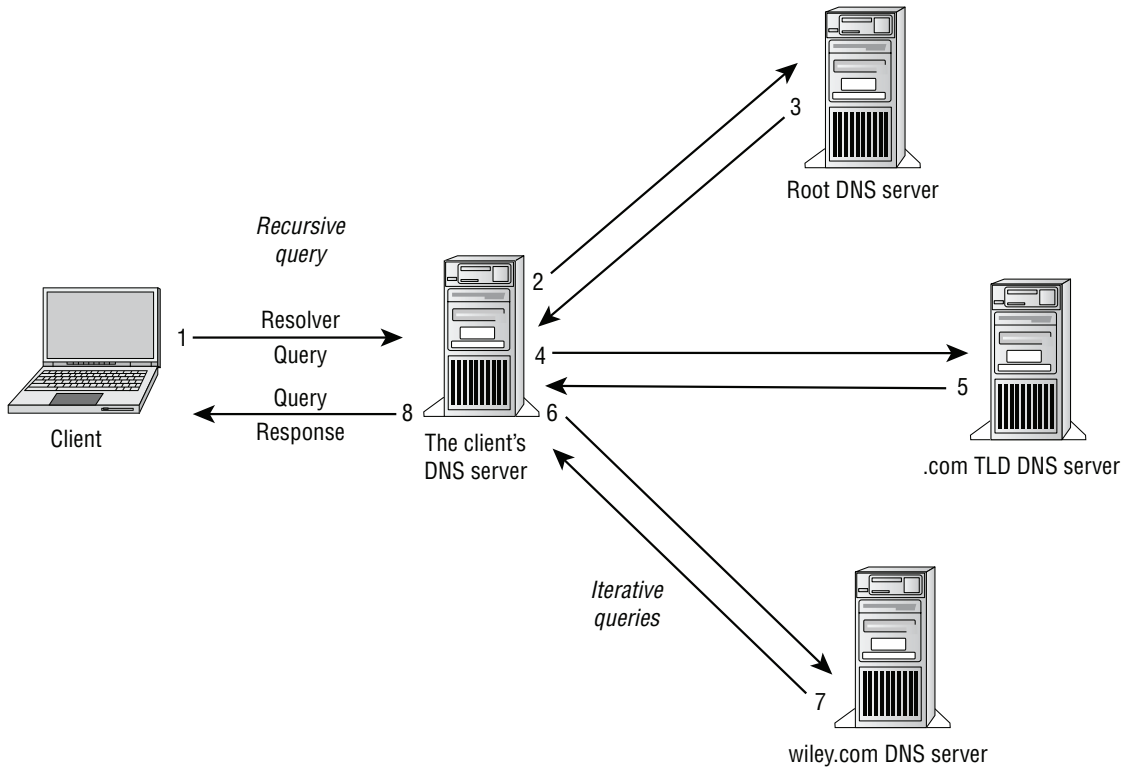

If you're curious, the list of root servers is maintained at `www.iana.org/domains/root/servers`, among other places.

After looking at Figure 8.4, you might be amazed that the Internet works at all. If nothing else, it should explain why, if you visit a website you've never visited before, it can sometimes take longer than normal to load (provided no one else who uses your DNS server has visited the site recently either). The next time you visit that site, it will probably appear faster, if the name resolution is still held in cache.

DNS servers for intranet use only can be located on the internal network (inside the network firewalls). If the DNS server is being used for Internet name resolution, it's most effective to place it in the screened subnet. Remember that if the network firewalls are blocking ports, DNS uses UDP or TCP port 53 so those will need to be opened.

DHCP

Dynamic Host Configuration Protocol (DHCP) was also briefly introduced in Volume 1 Chapter 7, and we cover it in more depth here. A *DHCP server* is configured to provide IP configuration information to clients automatically (dynamically), in what is called a *lease*. It's called that because the information is not permanently granted to the client computer, and the client must periodically request a renewed lease or a new lease. The following configuration information is typically provided in a lease:

- IP address
- Subnet mask
- Default gateway (the "door" to the outside world)
- DNS server address

DHCP as a server role is listed in Exam Objective 2.3. DHCP leases, scopes, reservations, and exclusions are part of Exam Objective 2.4. Understand and be able to explain how they work.

Scopes and Reservations

DHCP servers can provide a lot more than the items on the previous list, but these are the most common. The list of parameters that the DHCP server can provide is configured as part of a *scope*. DHCP servers can have one or more scopes to service clients on different subnets or network segments. The following items are typically included within the scope:

Address Pool This is the range of addresses that the server can give out to clients. For example, the pool may be configured to give out addresses in the range from 192.168.0.100 to 192.168.0.200. If all of the addresses are given out, then the server can't provide addresses to any new potential clients. If the network is using IPv4 addresses, the address pool configuration will also include the subnet mask.

Lease Durations IP addresses given out by the DHCP server are leased to clients, and the lease has an expiration time. Before the lease expires, the client (if it's online) will typically renegotiate to receive a new lease. If the lease expires, then the address becomes available to assign to another client. In situations where there are limited IP addresses but a lot of clients coming on and off the network frequently (say, a traveling sales force), shorter lease times are a good option. The downside is that it will generate a bit more network broadcast traffic, as we will explain in just a bit.

Address Reservations and Exclusions Some IP addresses can be reserved (it's appropriately named a DHCP *reservation*) for specific clients, based on the client's MAC address. This is particularly useful for devices that need to have a static IP address, such as printers, servers, and routers. In this way, all IP addresses can be accounted for and managed by the DHCP server, reducing potential confusion for administrators.

An *exclusion* is a range of addresses in the DHCP server's scope that it will not issue to clients. For example, an administrator can exclude a block of IP addresses and then manually assign those addresses to servers or printers that need static addresses.

Scope Options These provide extra configuration items outside of the IP address and subnet mask. The most common items are the address of the default gateway (the router) and DNS servers. Other items might include the addresses of servers providing other functions, such as time synchronization, NetBIOS name resolution, proxy server, telephony services, or the domain name (`whatever.com`) for the client to use.

Real World Scenario

TCP/IP Configuration Choices

When configuring TCP/IP on a network, you have three choices for assigning IP addresses: manually (also called *static IP addressing*), automatically using a DHCP server (also called *dynamic IP addressing*), and a hybrid approach.

The manual option takes the most work, and it really only works for smaller networks. The administrator needs to keep track of all the addresses that have been assigned so that the same address doesn't end up being accidentally assigned to multiple computers. Duplicate addresses will cause communication problems. Most administrators have better things to do than manage IP addresses manually.

DHCP is extremely convenient. The administrator sets up a scope and options and lets the server automatically manage all the IP addresses. For devices that need a static address, such as printers, servers, and routers, the administrator can configure address reservations based on their MAC addresses. That way, every time a specific printer comes online, it always receives the same address. This takes a bit of setup in the beginning, or when a new device is added, but it's worth it in the long run. The DHCP server is then the single point of management for all IP addresses on the network. That makes administration and troubleshooting much easier.

The hybrid option is a combination of manual assignments and DHCP. For example, devices that need static IP addresses can be assigned manually, whereas clients get their information from a DHCP server. In a situation like this, the administrator might set a DHCP client pool address range for 192.168.0.100 to 192.168.0.200 and then exclude addresses 192.168.0.1 to 192.168.0.99 for static devices. The problem with this approach is that it requires extra administrative effort to manage the static addresses and ensure that the same address isn't assigned multiple times. In addition, say another administrator looks at the DHCP server and sees the scope. That administrator might or might not know that addresses lower than .100 are being used for static assignments and could increase the scope to include those addresses. Doing so will cause the DHCP server to hand out addresses that could conflict with devices that have been manually configured. So hybrid is an option, but it's not recommended.

The best option is to use the DHCP server to automatically manage all IP addresses on the network.

How DHCP Works

DHCP clients need to be configured to obtain an IP address automatically. This is done by going into the network card's properties and then the TCP/IP properties, as you learned in Volume 1 Chapter 7.

When the client boots up, it will not have an IP address. To ask for one, it will send a DHCP DISCOVER broadcast out on the network. If a DHCP server is available to hear the broadcast, it will respond directly to the requesting client using the client's MAC address as the destination address. The process is shown in Figure 8.5.

FIGURE 8.5 The DHCP request process

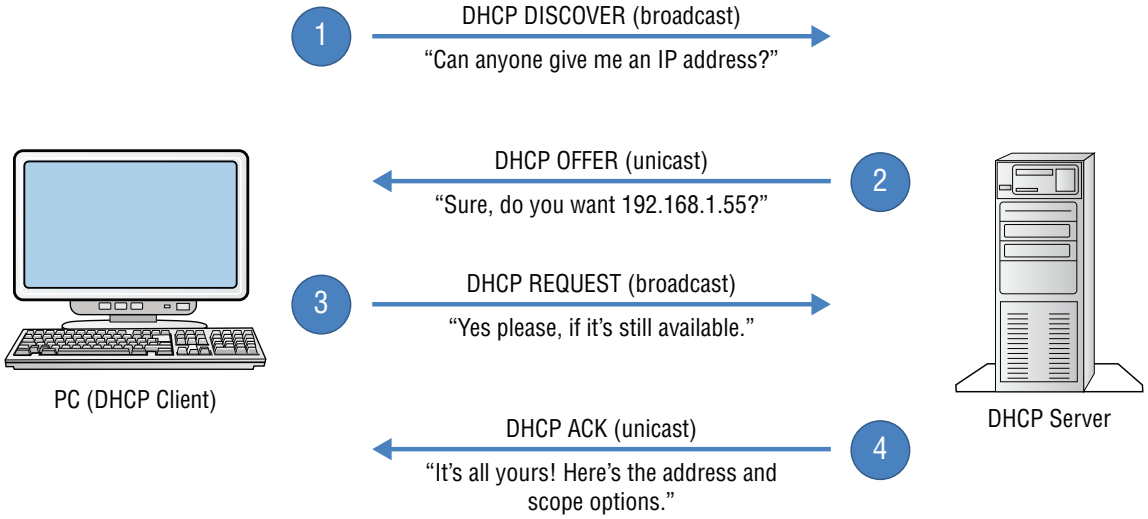

Notice that the DHCP DISCOVER and DHCP REQUEST messages are broadcasts, which means two important things. First, every computer on the network segment receives and needs to process the broadcast message. It's like snail mail that's addressed to "the current resident" at an address, and the computer is compelled to read it. Excessive broadcasts can dramatically slow network performance. Second, broadcasts do not go through routers. Thus, if the client and the DHCP server are on opposite sides of a router, there will be a problem. There are two resolutions. First, make the router the DHCP server. Second, install a *DHCP relay agent* on the subnet that doesn't have the DHCP server. It will be configured with the address of the DHCP server, and it will forward the request directly to the DHCP server on behalf of the client.

DHCP servers should be located on the internal network. If the network has clients that are connecting via remote access, then a device with DHCP capabilities (such as the Routing and Remote Access Service [RRAS]) can be placed in the screened subnet. DHCP uses UDP ports 67 and 68.

Fileshare or File Server

A *fileshare* or *file server* provides a central repository for users to store, manage, and access files on the network. There are a few distinct advantages to using file servers:

- Ease of access to files for collaboration
- Centralized security management
- Centralized backups

Fileshares are part of Exam Objective 2.3. Understand what they are and how they are used.

Fileshares come in a variety of shapes and sizes. Some are as basic as Windows- or Linux-based servers with a large amount of internal hard disk storage space. Networks can also use *network-attached storage (NAS)* devices, which are stand-alone units that contain hard drives, come with their own file management software, and connect directly to the network. If a company has extravagant data storage needs, it can implement a *storage area network (SAN)*. A SAN is basically a network segment, or collection of servers, that exists solely to store, manage, and access data quickly.

Since the point of a fileshare is to store data, it's pretty important to ensure that it has ample disk space. Some dedicated file servers also have banks of multiple optical drives for extra storage or for letting users access files from optical media. Processing power and network bandwidth can also be important to manage file requests and deliver them in a timely manner.

If fileshares are being managed on the corporate network, they will almost always be on the internal network. There might be situations where a fileshare is also an FTP server, in which case the server should be on the screened subnet. In those cases, however, ensure that the server does not contain highly sensitive information or other data that the company doesn't want to lose. It's becoming increasingly common for fileshares to be stored in the cloud (which we will cover in Volume 1 Chapter 11, "Virtualization and Cloud Computing"). In instances such as that, physical security will be the responsibility of the cloud provider. Client computers may be required to use a VPN to connect to it.

Print Server

Print servers are much like file servers, except, of course, they make printers available to users. In fact, file servers and print servers are combined so often that you will see a lot of publications or tools refer to *file and print servers* as if they were their own category.

Print servers are part of Exam Objective 2.3. Be able to summarize the role they play on a network.

On its own, a *print server* makes printers available to clients over the network and accepts print requests from those clients. A print server can be a physical server like a Windows- or Linux-based server, a small stand-alone device attached to a printer (or several printers), or even a server module built into the printer itself. Print servers handle the following important functions:

- Making printers available on the network
- Accepting print requests
- Managing print requests (in the print queue)
- In some cases, processing and storing print jobs

Figure 8.6 shows a simple stand-alone print server. It has an RJ45 network connection and four USB ports to connect printers. Wireless print servers are easy to find as well. Some are as inexpensive as $40–$50, whereas others can cost several hundred dollars.

FIGURE 8.6 A D-Link print server

Although the specific functionality will vary by print server, most of the time administrators will be able to manage security, time restrictions, and other options, including who can access the printer, if the server processes the files, and if the print jobs are saved after printing. An example is shown in Figure 8.7—you may recall these from Volume 1 Chapter 5, "Printers and Multifunction Devices." Print servers should be located on the internal network.

Mail Server

Email is critical for communication, and mail servers are responsible for sending, receiving, and managing email. To be a *mail server*, the computer must be running a specialized email server package. Some popular ones are Postfix, Exim, and Microsoft Exchange Online, although there are dozens of others on the market.

FIGURE 8.7 Printer management options

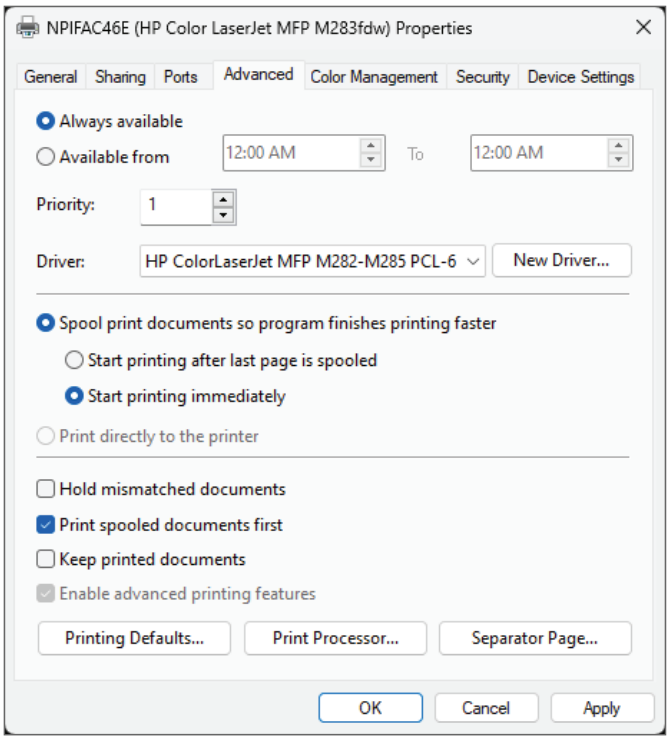

Exam Objective 2.3 includes mail servers as a subobjective. Understand the service they provide.

Clients access the mail server by using an email client installed on their systems. The most common corporate email client is Microsoft Outlook, but Gmail, Apple Mail, Thunderbird, and others are also used. Mobile and Internet email clients (which are more popular than their corporate cousins) are dominated by Gmail for Android clients and Apple Mail for the iPhone and iPad. Other options include Outlook, Mozilla Thunderbird, Yahoo! Mail, and you'll still see an occasional AOL mail diehard around too. (In many circles it's perfectly acceptable to make fun of people who use AOL email addresses.)

In addition to sending and receiving email, mail servers often have antispam software built into them as well as the ability to encrypt and decrypt messages. Email servers are most often located in the screened subnet. Table 8.2 reviews the most important protocols for sending and receiving email that you learned about in Volume 1 Chapter 7. Firewalls need to allow traffic on these ports through for email to work properly.

TABLE 8.2 Important email protocols

Protocol	Port	Purpose
SMTP	25	Sending email and transferring email between mail servers.
POP3	110	Receiving email.
IMAP4	143	Receiving email. It's newer and has more features than POP3.

SMTP is a push (send) protocol, whereas POP3 and IMAP4 are pull (receive) protocols.

Syslog

Network administrators need to know what's happening on their network at all times. The challenge is that there may be hundreds or thousands of devices on the network, with thousands of users accessing resources locally and remotely. Keeping track of who is logging in where, what resources users are accessing, who is visiting the web server, the status of the router, the printer's online status, and innumerable other events could be an administrative nightmare. Fortunately, *syslog* is available to help manage it all.

Exam Objective 2.3 includes syslog as a subobjective. Understand how syslog works and the roles that clients and servers play in its operation. Syslog works as a client-server model, where the clients generate messages based on the triggering of certain conditions, such as a login event or an error with a device, and send them to a centralized logging server, also known as the *syslog server*. Consequently, the term *syslog* can be applied to a standard or system for event monitoring, the protocol, or the actual server that collects the logged messages.

Syslog got its start in the UNIX world and is used extensively with Linux-based networking systems and devices. Microsoft operating systems don't natively support syslog—Windows comes with its own event logger called Event Viewer, which we cover in Volume 2 Chapter 2, "Windows Configuration"—but it's easy to find packages that let Windows servers participate in a syslog environment. Let's take a look at clients and servers in a syslog system.

Syslog Clients

Many different types of devices, such as servers, routers, and printers, support syslog as a client across a wide variety of operating systems. The primary job of the client (in syslog terms) is to send a message to the syslog server if certain conditions are met. For example,

an authentication server can send a message whenever there is a successful or failed login attempt, or a router can send the status of its used or available bandwidth.

Messages have the following three components:

A Facility Code The *facility code* is a number between 0 and 23 that identifies the type of device sending the message. For example, facility code 0 is for operating system kernel messages; code 2 is used by email servers; code 4 is for security messages; and code 6 is used by printers. Don't worry about memorizing the facility codes. Just know that they tell the server where the message came from.

A Severity Level The *severity level* lets the administrator know how urgent the issue is. Table 8.3 illustrates the syslog severity levels. The lower the level, the more urgent the message is. Facility codes and error levels can be helpful when sorting events on a syslog server, and they can also be used to send the administrator an alert if something catastrophic happens. Facility codes and severity levels are required message components.

A Text Description The final portion of the message is the text description of the message itself. The description may be in easy-to-read language, or it could be nonsensical ranting. It really depends on the device developer and what they program it to be. There is no standardized format for the description, so it can come in any number of formats. Most messages do come with an IP address or device name included, but the lack of standardization can make understanding messages a fun and entertaining challenge.

TABLE 8.3 Syslog severity levels

Level	Severity	Description
0	Emergency	A panic condition when the system is unusable
1	Alert	Immediate action needed
2	Critical	Major errors in the system
3	Error	"Normal" error conditions
4	Warning	Warning conditions, usually not as urgent as errors
5	Notice	Normal operation but a condition has been met
6	Information	Provides general information
7	Debug	Information used to help debug programs

Syslog Servers

The syslog server's job is to collect and store messages. Most syslog servers are made up of three components: the listener, a database, and management and filtering software.

Syslog servers listen on UDP port 514 by default. Remember that UDP is a connectionless protocol, so the delivery of packets is not guaranteed. The default implementation of syslog is also not secure. However, it can be secured by running syslog over Transport Layer Security (TLS) and TCP port 6514. Regardless of whether it's secured or not, always place the syslog server behind the firewall and on the internal network.

Even on small networks, devices can generate huge numbers of syslog messages. Therefore, most syslog implementations store messages in a database for easier retrieval and analysis.

Finally, most syslog servers will have management software that is used to view messages. The software should also have the ability to send the administrator a console message or text (or email) if a critical error is logged. Dozens of syslog packages are available. Some popular packages are Kiwi Syslog Server NG by SolarWinds (shown in Figure 8.8), Splunk, syslog-ng, and Syslog Watcher.

FIGURE 8.8 Kiwi Syslog

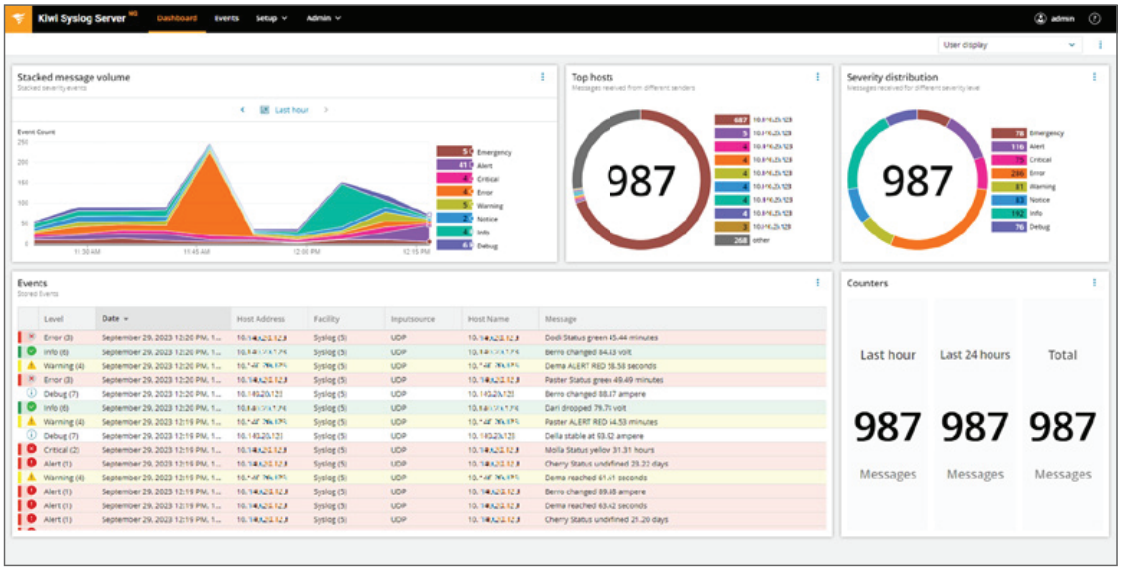

Web Server

Whenever you visit a web page, you are making a connection from your device (the client) to a *web server*. To be more specific, a connection is requested by your Internet software (generally, a web browser) using the *Hypertext Transfer Protocol Secure (HTTPS)* protocol of the TCP/IP protocol suite. Your client needs to know the IP address of the web server, and it will make the request on port 443.

Web servers are part of Exam Objective 2.3. Know how they work and what ports they use.

The web server itself is configured with web hosting software, which listens for inbound requests on port 443. The most common web server platforms are the open source Nginx and Apache, and LiteSpeed and Microsoft's Internet Information Services (IIS). Web servers provide content on request, which can include text, images, and videos, and they can also do things like run scripts to open additional functions, such as processing credit card transactions and querying databases.

Web servers often function as download servers, too, meaning that they also use FTP. FTP works on TCP ports 20 and 21.

Individuals or independent companies can manage web servers, but more often than not an Internet service provider or web hosting company that manages hundreds or thousands of websites manages them. In fact, one web server can be configured to manage dozens of smaller websites using the same IP address, provided that it has sufficient resources to handle the traffic. On the flip side, huge sites, such as Amazon and Google, are actually made up of multiple web servers acting as one site. It's estimated that Google has over 900,000 servers, and Microsoft claims to have over 1 million servers!

If a company wants to host its own web server, the best place for it is in the screened subnet. This configuration provides ease of access (after all, you *want* people to hit your web server) and the best security. The firewall can be configured to allow inbound port 443 requests to the screened subnet but not to allow inbound requests on those ports to make it to the internal corporate network.

Contrast this to a situation where the web server is on the internal network. The firewall then has to let inbound port 443 connections through to the internal network so that Internet-based clients can get to the web server. However, that also means that inbound requests on port 443 can be sent to all internal computers, including non-web servers and even client computers. Hackers could then potentially take advantage of exploits using port 443 to attempt to gain illegitimate access to the network.

Authentication, Authorization, and Accounting (AAA)

The ultimate goal of a security system is to protect resources by keeping the bad people out and letting the good people in. It would be really easy to configure a system such that no one could access anything, and it would be equally simple to let everyone have open access. The first extreme defeats the purpose of having a network, and the second is just begging for trouble. The challenge then is to find a happy medium, where resources are available to those who should have them and nobody else.

AAA servers are part of Exam Objective 2.3. They are critically important to any network, and you will be tested on them as well. Be sure to understand the differences between the three As and how they work together to secure network resources.

In information security, there's a framework for access control known as *triple A*, meaning *authentication, authorization, and accounting (AAA)*. Occasionally auditing is added to the mix, making it *quad A*. And even further, nonrepudiation, or the assurance that something can't be denied by someone, is also sometimes lumped in. Regardless, triple A is the umbrella term for describing systems of access control. AAA servers are gatekeepers and critical components to network security, and they can be implemented on a dedicated server machine, wireless router or access point, Ethernet switch, or a remote access server.

A common term that you will hear in the Windows Server world is *domain controller*, which is a centralized authentication server as part of its Active Directory (AD) structure. Other types of servers that handle all aspects of AAA are *Routing and Remote Access Service (RRAS), Remote Authentication Dial-In User Service (RADIUS)*, *Terminal Access Controller Access-Control System Plus (TACACS+)*, and *Kerberos*. Authentication servers may be standalone (e.g., a "Kerberos server"), or the authentication service may be built into a more well-known OS. For example, Windows Server uses Kerberos.

The AAA process will differ slightly between servers, but generally what happens is the user (or computer) trying to access the network presents credentials. If the credentials are deemed appropriate, the authentication server issues the user a security code or a ticket that grants them access to resources. When the owner of the security code or ticket tries to access a resource, authorization comes into play. And finally, accounting tracks all of it. In the following sections, we will describe the principles of authentication, authorization, and accounting.

Authentication

To implement security, it's imperative to understand who or what is accessing resources on a computer or network. User *authentication* happens when the system being logged into validates that the user has proper credentials. Essentially, the authentication server asks the question, "Who are you?" Oftentimes, this is as simple as entering a username and password, but it could be more complex. There are two categories of authentication:

Single-Factor Authentication The simplest form of authentication is *single-factor authentication (SFA)*. A single-factor system requires only one piece of information beyond the username to allow access. Most often, this is a password. Single-factor authentication is quite common, but it's not the most secure method out there.

Multifactor Authentication To increase security, your computer or network might require *multifactor authentication (MFA)*, which as the name implies requires multiple pieces of information for you to log in. Generally speaking, in addition to a username, multifactor authentication requires you to provide two or more pieces of information (called authentication factors) out of these four categories: something you know, something you have, something you are, or somewhere you are.

You may also hear multifactor authentication referred to by the specific number of factors required, such as two-factor authentication or three-factor authentication.

As mentioned already, most authentication systems require just a password, which is an example of something you know. If you forget your password, a website might ask you to provide answers to security questions that you selected when you registered. These are questions such as the name of your elementary school, father's middle name, street you grew up on, first car, favorite food, musical artist, and so forth.

One-time passwords can be generated by sites to give you a limited time window to log in. These are far more secure than a standard password because they are valid for only a short amount of time, usually 30 minutes or less. The password will be sent to you via text or email or possibly a phone call.

Something you have can be one of a few different things, such as a smartcard or a security token. A smartcard is a plastic card, similar in dimensions to a credit card, which contains a microchip that a card reader can scan, such as on a security system. Smartcards often double as employee badges, enabling employees to access employee-only areas of a building or to use elevators that go to restricted areas, or as credit cards.

Smartcards can also be used to allow or prevent computer access. For example, a PC may have a card reader on it through which the employee has to swipe the card, or that reads the card's chip automatically when the card comes into its vicinity. Or, they're combined with a PIN or used as an add-on to a standard login system to give an additional layer of security verification. For someone to gain unauthorized access, they have to know a user's ID and password (or PIN) and also steal their smartcard. That makes it much more difficult to be a thief!

A security token, like the one shown in Figure 8.9, displays an access code that changes about every 30 seconds. When received, it's synchronized with your user account, and the algorithm that controls the code change is known by the token as well as your authentication system. When you log in, you need your username and password, along with the code on the token.

FIGURE 8.9 RSA SecurID

Security tokens can be software-based as well. A token may be embedded in a security file unique to your computer, or your network may use a program that generates a security token much like the hardware token does. Figure 8.10 shows an example of PingID, which works on computers and mobile devices. This type of token saves you from having to carry around yet another gadget.

FIGURE 8.10 PingID

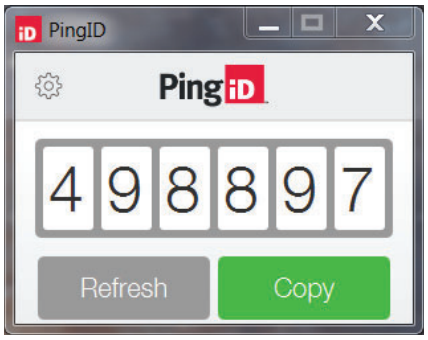

A system might also require users to log in from a specific location. For example, perhaps users are allowed to log in only if they are on the internal corporate network. Or, maybe they are allowed to connect from their home office. In that case, the security system would know a range of IP addresses to allow in based on the block of addresses allocated to the user's ISP. This is an example of somewhere you are.

Finally, the system could require something totally unique to you (something you are) to enable authentication. These characteristics are usually assessed via biometric devices, which authenticate users by scanning for one or more physical traits. Common types include fingerprint recognition, facial recognition, and retina scanning.

Authorization

Once it's determined who the user is, the next step in access control is determining what the user can do. This is called *authorization*. Users are allowed to perform only specific tasks on specific objects based on what they are authorized to do. Most computers grant access based on a system of *permissions*, which are groups of privileges. For example, a user might be able to make changes to one file, whereas they are only allowed to open and read another.

One of the key foundations of an authorization system is the principle of *least privilege*. This states that users should be granted only the least amount of access required to perform their jobs, and no more. This principle applies to computers, files, databases, and all other available resources.

Permissions are important! We cover types of permissions (such as file system vs. share permissions) in Volume 2 Chapter 6, "Securing Operating Systems," and restricting user permissions in Volume 2 Chapter 7, "Troubleshooting Operating Systems and Security."

Accounting

After users have been authenticated and authorized, it's time to think about tracking what the users did with their access. This is where *accounting* comes in. The principle of

accounting seeks to keep a record of who accessed what and when, and the actions they performed.

The most common method of tracking user actions is through the use of logs. Nearly all operating systems have built-in logs that track various actions. For example, Windows-based systems contain Windows Logs, which are part of Event Viewer. As noted earlier in this chapter, we will cover Event Viewer in more depth in Volume 2 Chapter 2, but let's take a quick look here. To open Event Viewer, type **Event** into the Windows search bar and press Enter. Windows has logs that track application events, security events, and system events. Figure 8.11 shows the Security log. In an environment where multiple users log in, those log-ins will be shown here.

FIGURE 8.11 Security log in Event Viewer

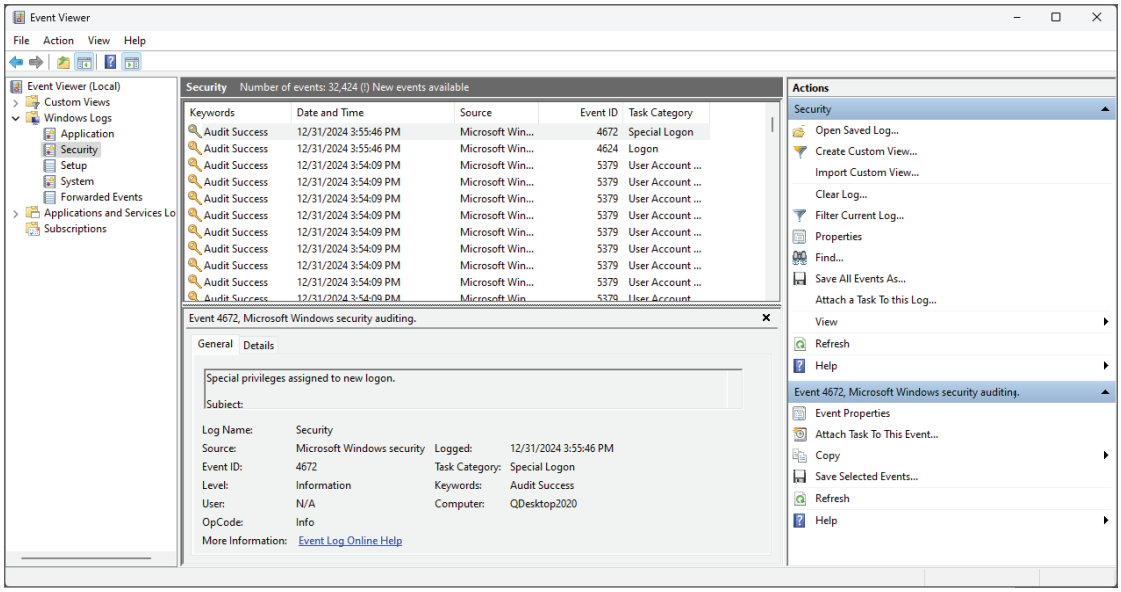

Another action that is frequently tracked is web browsing history. Web browsers retain a historical account of the sites that have been visited. To see viewing history in Microsoft Edge, click the Settings and More menu (the three horizontal dots near the upper-right corner) and then History, as shown in Figure 8.12. There's an option to clear the history as well. Note that this action clears it from the browser, but it won't clear it from any servers (such as a proxy server) that caches web requests. To view the history in Chrome, click the More menu (the three vertical dots) and then History, or just open Chrome and press Ctrl+H.

Database Servers

As their name suggests, *database servers* store databases, which are organized collections of information or data. They run specialized software called a *database management system*

(DBMS). Traditional examples include Oracle, Microsoft SQL Server, and MySQL, and there are dozens of others commercially available. Cloud-based database solutions are popular and growing, including Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform.

FIGURE 8.12 Microsoft Edge site-viewing history

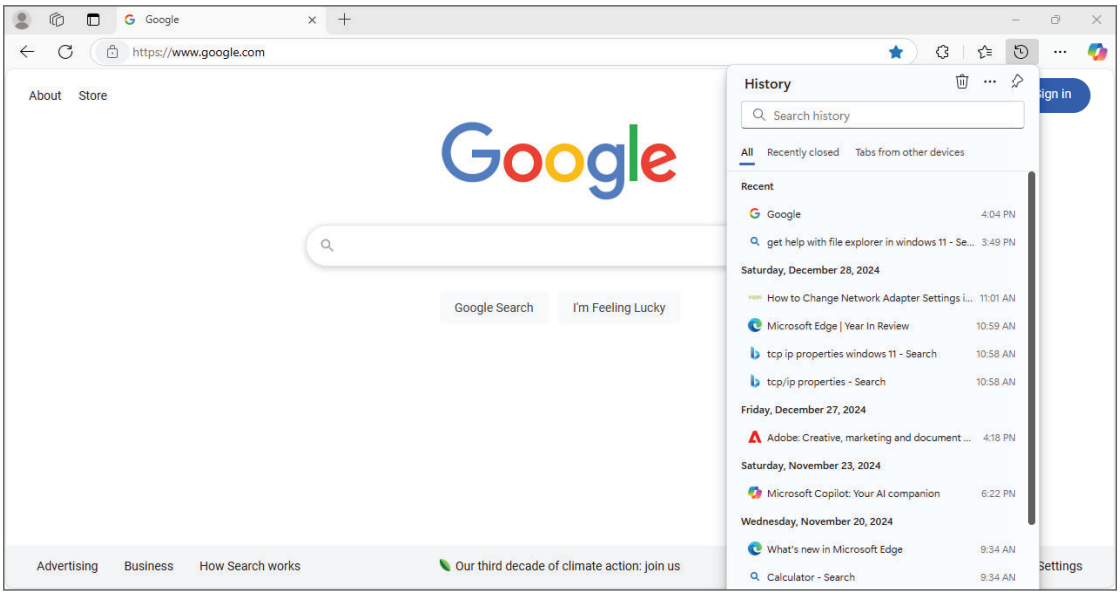

Database servers are part of Exam Objective 2.3. Understand what they do and be able to name a few different database software packages.

Databases are used in a variety of business settings. Banks, online retailers, government agencies, nonprofits, and even bloggers use databases to store data that's important to them or their customers. There isn't a one-size-fits-all solution, and it might take some investigation to find the right solution for the situation. The differences are in how they store it and, more importantly, how a user extracts data from the database. Regardless, though, the concepts of storage and access are pretty consistent across platforms.

When most people think of a database, they often think about a list with columns similar to a Microsoft Excel worksheet, like the one shown in Figure 8.13. Each column has a certain type of data, such as an identifier, name, phone number, or email address. Other sheets might have data such as a date, part number, amount, product description, and so on. One difference is that in a database, the structure shown in Figure 8.13 would be called a *table* (instead of a *worksheet*). Otherwise, a basic database with a flat structure looks a lot like something you could see in Excel.

FIGURE 8.13 Simple data in a worksheet

	A	B	C	D	E	F	G
1	**ID**	**Last**	**First**	**Phone**	**Email**		
2	C1	Smith	Joe	(123) 456-7890	joesmith@fakeemail.com		
3	C2	Larson	Jane	(123) 456-7891	janesmith@notreal.com		
4	C3	Hernandez	Maria	(111) 234-5678	mh123@notreal.com		
5	C4	Calderon	Juan	(112) 345-6789	juanc@fakeemail.com		
6	C5	White	Michael	(113) 456-7890	mcwhite@notreal.com		
7	C6	Jefferson	Mary	(124) 123-4567	maryj@fakeemail.com		
8							

The Excel example shown in Figure 8.13 is what's known as a *flat file*. It's called that because the data has two dimensions, rows and columns. Databases don't need to have a flat structure—and in fact most don't—which is one of their advantages over spreadsheet software. This allows databases to manage far more complicated relationships between data points than spreadsheets can.

Databases, and therefore database servers, provide the following benefits over flat files:

Multiple Concurrent Users If multiple concurrent users need to access data, a database is needed. Flat files don't do well when several people try to access them at once. Online collaboration programs have gotten better, but databases are still the gold standard when it comes to concurrent use. This is because of how record locks are employed by flat files versus databases.

When a user tries to write to a flat file, the whole file is locked until the function executes. Databases generally lock only on a record level (or individual line of data), so multiple people can actively use the database at one time.

Some data collaboration software packages will let multiple users "edit" a spreadsheet at the same time. This is a bit of sleight of hand. The collaboration software provides for the functionality by essentially holding multiple copies of the spreadsheet, tracking the changes, and then writing the changes based on a time stamp. The spreadsheet software isn't aware that multiple people are editing it.

Scalability Database servers are extremely scalable, allowing access to hundreds or tens of thousands of users at a time. Enterprise versions of database software also allow databases to stretch between on-site servers and the cloud, with the most recent data being available on-site.

Speed Database software is very memory intensive, so having plenty of RAM in a database server is critical. Also, speed can be increased by placing databases and their log files on SSDs rather than HDDs. Admittedly, finding large enough SSDs can be an issue for large databases, but the speed difference helps.

Variety of Data Spreadsheets are great at storing numbers and text, but not so great at storing images or even entire files within themselves. Databases can store virtually any type of data and make it available quickly when searching for it.

Number of Records Excel worksheets are limited to no more than 1,048,576 rows and 16,384 columns. In addition, one cell can't contain more than 32,767 characters. While this might sound like a lot, think about large organizations. Does Amazon serve more than one million customers? Probably so—just today. How about the U.S. Internal Revenue Service? The IRS needs to keep track of the Social Security numbers of roughly 350 million living Americans and many more deceased ones too. A spreadsheet just isn't going to cut it.

Data Persistence Users expect their data to be stored indefinitely, and this is what is meant by *data persistence*. Spreadsheets and databases both write data to hard drives, so by definition, both possess this trait. However, imagine a situation where you are entering data and the computer loses power. If you were editing a spreadsheet, all changes since the last file save are gone. With a database, all changes are essentially saved after each cell of data is edited. Furthermore, databases log changes in separate log files. If something happens to corrupt the data, the log files can be used to re-create the missing or corrupted data. Therefore, databases have much better data persistence than do spreadsheets.

Data Availability Flat files are meant to be stored on a local computer. This limits the number of people who can access them. Databases can be stored on a local computer, called *offline storage*, but they can also be stored on a server or in the cloud to enable broad access, called *online storage*. Each of these options enhances *data availability*, which means data is available to all users at any time of day and regardless of location.

Security Databases provide much more granular security than do spreadsheets. With a spreadsheet, either a user can open the file and read all of its contents or the user can't. By using databases, users might be able to access specific columns or rows of data but not see anything else. For example, a customer might be able to look up all of their orders in the database, but they won't be able to see orders from other customers.

Databases managed on a local network should almost certainly be placed behind network firewalls. Database software can work hand-in-hand with firewalls to manage access if external users need to get to the data.

Types of Databases

This is likely beyond the scope of the A+ exam but may be interesting to you anyway. There are two major classifications of databases in use today: relational databases and non-relational databases.

A relational database has a rigid structure, much like the Excel example we shared in this section. It contains tables with columns and rows to store data. The rules of the database, such as what type of data can be stored and its structure, is called the schema.

Relational databases are great for storing structured data, but most data in the world—80 percent by some professional estimates—is nonstructured. This includes things like pictures, videos, web pages, emails, documents, texts, and social media posts. This type of data is best stored in a non-relational database. Examples include key/value databases and document databases. "Big Data," which is a popular business buzzword these days, requires the use of non-relational databases.

If you find databases interesting, there are many career opportunities available in database management and data analytics. And as you might expect, there are certifications as well. They include CompTIA's DataSys+ as well as Microsoft, Google, Oracle, and other options.

Network Time Protocol

A *Network Time Protocol (NTP) server* is one that's designated as the reliable source of time for devices on a network. It allows all devices to synchronize their time with the server and each other. This might not sound like a big deal, but when it comes to concepts like AAA, having accurate and aligned time stamps on security events becomes crucial.

NTP servers are part of Exam Objective 2.3. Be able to explain what they do on a network.

Most networks don't need to run their own internal NTP server. There are hundreds of public NTP servers available, such as `time.google.com`, `time.aws.com`, and `time.nst.gov`. Windows servers and clients can synchronize using the Windows Time Service and the `time.windows.com` server. Almost all regions and countries across the globe have time servers as well, such as `time.fu-berlin.de` in Germany.

Some organizations may have a need for internal NTP servers. Reasons could include compliance regulations, offline functionality, synchronization of automated tasks, or enhanced security because there is no reliance upon public servers. (A local NTP server *can* be set up to query public servers but does not need to.) In cases like those, the NTP server should be behind all firewalls. NTP operates on UDP port 123.

Internet Appliances

The definition of an *Internet appliance* is a device that makes it easy to access the Internet. Taking a slightly broader view, Internet appliances can also help users safely access the Internet by protecting against some of the dangers that lurk there. The CompTIA A+ 220-1201 exam objectives list four items under Internet appliances: spam gateways, unified threat management (UTM), load balancers, and proxy servers. Let's take a look at each one.

Spam Gateways

Spam email is pervasive today. If you use email, it's likely you get spam, and a lot of it. What might be hard to believe is that you would get significantly more if there weren't antispam devices protecting you. A *spam gateway* is an appliance—most likely a software installation

or virtual appliance—that blocks malicious emails from entering a network. They go by other names as well, such as antispam gateways (which sound more appropriate), spam blockers, and email gateways.

Spam gateways are part of Exam Objective 2.3. Understand their role on a network.

Antispam gateways can be located in two places: on the cloud or on an internal network, meaning internal to where the firewall is placed. The intent is that emails inbound to a corporate email server will first go through the gateway. If the gateway checks the email and verifies it's not spam, it passes the mail through to the server. Flagged emails get sent to a spam folder, quarantined, or deleted.

Emails that contain certain keywords, have malicious links, or come from domains that are known to send spam are most likely to get flagged. On occasion, legitimate emails will get flagged too, but these appliances get far more right than they get wrong.

Some spam gateways will also handle outbound emails. Hopefully your users aren't spammers, but if they are, the gateway will put an end to it. In most cases, a company sending spam email is accidental. A hacker could be sending emails from your company's domain by spoofing it or by other tricks which might not be immediately visible. A good spam gateway will block these outbound messages and notify the administrator.

Unified Threat Management (UTM)

The Internet is a wondrous place, but it's a scary one as well. It seems like for every video of puppies or kittens doing cute things, there are 10 hackers lurking in dark corners trying to steal identities or crash servers. It's an unfortunate reality of the Internet age. Software and hardware solutions have sprung up in response to various types of threats and managing all of them can be a challenge. For example, a network needs a firewall, antimalware and antispam software, and perhaps content filtering and intrusion prevention system (IPS) devices as well. It's a lot to deal with.

The goal of *unified threat management (UTM)* is to centralize security management, allowing administrators to manage all their security-related hardware and software through a single device or interface. For administrators, having a single management point greatly reduces administration difficulties. The downside is that it introduces a single point of failure. If all network security is managed through one device, a device failure could be problematic.

Exam Objective 2.3 lists UTM as a subobjective. Understand how it works to coordinate and centralize security management.

UTM is generally implemented as a stand-alone device (or series of devices) on a network, and it will replace the traditional firewall. A UTM device can generally provide the following types of services:

- Packet filtering and inspection, like a firewall
- Intrusion detection and prevention services

- Gateway antimalware
- Spam blocking
- Malicious website blocking (either prohibited or nefarious content)
- Application control

UTM has become quite popular in the last several years. Many in the industry see it as the next generation of firewalls.

Detecting and Preventing Network Intrusions

When marketing unified threat management (UTM) systems, many vendors will talk about how their system combines an *intrusion detection system (IDS)* or *intrusion prevention system (IPS)* with other features. Both devices are closely related, and both devices monitor network traffic and look for suspicious activity that might be the sign of a network-based attack. You can think of them as being somewhat analogous to antivirus (AV) programs. AV programs examine individual files looking for telltale signs of malicious content (called a *signature*). IDSs and IPSs look for signatures as well, but in network traffic patterns.

Both IDSs and IPSs are different from firewalls. Firewalls have specific sets of rules that allow or deny packets to enter a network, based on criteria such as the packet's origination or destination address, the protocol being used, or the source or destination port. Their primary function is to block malicious traffic from entering the network in the first place. Firewalls may also have IDS or IPS software built in, or IDS and IPS devices can be stand-alone hardware devices.

An IDS is a passive device. It watches network traffic, and it can detect anomalies that might represent an attack. For example, if an attacker were to try to flood a network with traffic on a specific port, the IDS would sense that the additional traffic on that port was unusual. Then the IDS would log the anomaly and send an alert to an administrator. Note that it does nothing to prevent the attack; it simply logs relevant information pertaining to the attack and sends an alert.

By contrast, an IPS is an active device. It too monitors network traffic, but when it detects an anomaly, it can take actions to attempt to stop the attack. For example, if it senses suspicious inbound traffic on a specific IP port, it can shut down the port, block the sender, or reset the TCP connection. If the malicious stream of data is intended for one computer, it can prevent the attacking host from communicating with the intended victim computer or prevent the attacker from communicating with any computer. The specific actions it can take will depend on the device.

Both types of devices come in network-based varieties (NIDS and NIPS) and host-based varieties (HIDS and HIPS). As you might expect based on their names, the network-based versions are designed to protect multiple systems, whereas the host-based ones protect only one computer.

Load Balancing

Imagine you want to do some online shopping. You open your browser, type **amazon.com** into the address bar, and the site appears. You've made a connection to the Amazon server, right? But is there only one Amazon server? Considering the millions of transactions Amazon completes each day, that seems highly unlikely. In fact it's not the case. Amazon has hundreds of web servers, each of them capable of fulfilling the same tasks to make your shopping experience as easy as possible. Each server helps balance out the work for the website, which is called *load balancing*.

Load balancing appliances are part of Exam Objective 2.3. Understand the benefits of load balancing and be able to explain a few examples of how it works.

Load-balancing technology can be implemented with local hardware or on the cloud. If implemented on a local network, a hardware device, conveniently named a load balancer, essentially acts like the web server to the outside world. When a user visits the website, the load balancer sends the request to one of many real web servers to fulfill the request. Cloud implementations have made load balancing easier to configure and expand, since the servers can be virtual instead of physical.

Common Load Balancing Configurations

We already shared one example of load balancing with an online retailer. In that example, all servers are identical (or very close to identical) and perform the same tasks. Two other common load-balancing configurations are cross-region and content-based.

In a cross-region setup, all servers likely provide access to the same types of content, much like in our Amazon example. The big feature with this setup is that there are servers local to each region—proximity to the users will help speed up network performance. For example, say that a company has a geo-redundant cloud and users in North America, Asia, and Europe. When a request comes in, the load balancer senses the incoming IP address and routes the request to a server in that region. This is illustrated in Figure 8.14. If all servers in that region are too busy with other requests, then it might be sent to another region for processing.

Another common way to load-balance is to split up banks of servers to handle specific types of requests. For example, one group of servers could handle web requests, while a second set hosts streaming video and a third set manages downloads. This type of load balancing is called content-based load balancing and is shown in Figure 8.15.

Load Balancing Benefits

Load balancing has performance benefits for high-traffic networks and heavily used applications. Scalability and reliability are important benefits as well. Let's look at a few examples of each.

FIGURE 8.14 Cross-region load balancing

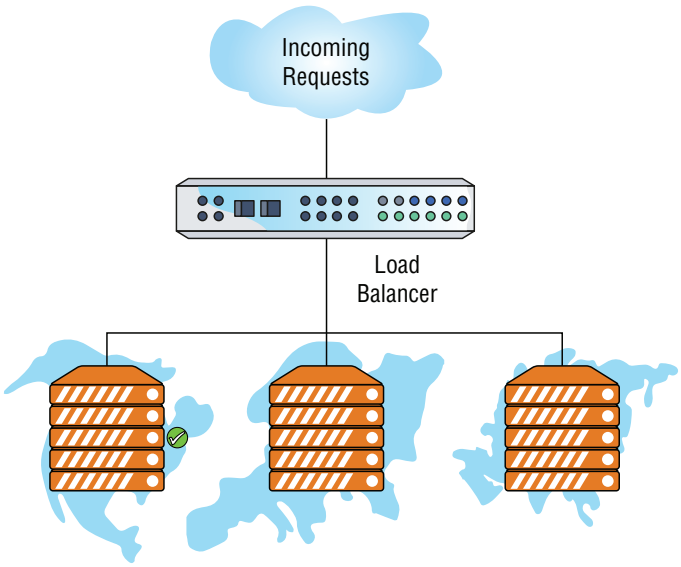

Performance The Amazon example we used earlier is probably the best example of this, but not all companies provide massive online shopping services. Smaller companies can benefit from performance enhancements as well. Servers that are specialized to handle a specific content type are often more efficient than multipurpose ones. And, the global load balancing example can be applied to distributed sites within a country as well.

Scalability If using cloud-based load balancing, services are easily scalable. For example, let's say that your company sells products online, and has two servers dedicated to that task. For the majority of the year, the two servers can handle the traffic without a problem. On the busiest shopping day of the year—Cyber Monday—those servers are overwhelmed. With cloud-based load balancing, traffic spikes can be handled by quickly provisioning additional virtual servers to handle the traffic. When the capacity is no longer required, the virtual servers are turned off.

Reliability Imagine a company that uses a business-critical application for remote salespeople. What happens if the server hosting that application crashes? It wouldn't be good.

With load balancing, different servers can host the application, even in different regions. Perhaps a hurricane wipes out the data center in Florida. The load balancer can direct users to other data centers in different regions, and the business can continue to generate revenue.

FIGURE 8.15 Content-based load balancing

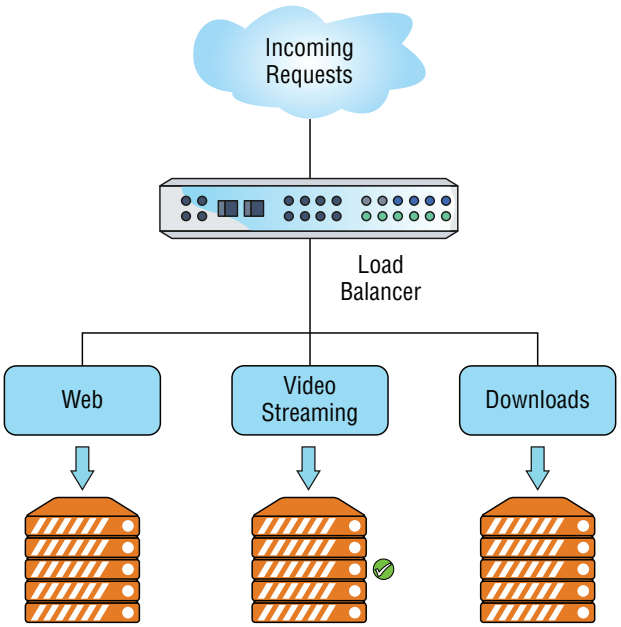

Proxy Server

A *proxy server* makes requests for resources on behalf of a client. The most common one that you will see is a web proxy, but you might run into a caching proxy as well. Exercise 8.1 shows you where to configure your computer to use a web proxy server in Windows 11.

EXERCISE 8.1

Configuring Windows 11 to Use a Proxy Server

1. Click Start, then Settings (it looks like a gear). That will open the Settings app.
2. On the left side, click Network & Internet.
3. In the list of Network & Internet options, click the Proxy tile. You will get a screen like the one shown in Figure 8.16.

FIGURE 8.16 Proxy settings

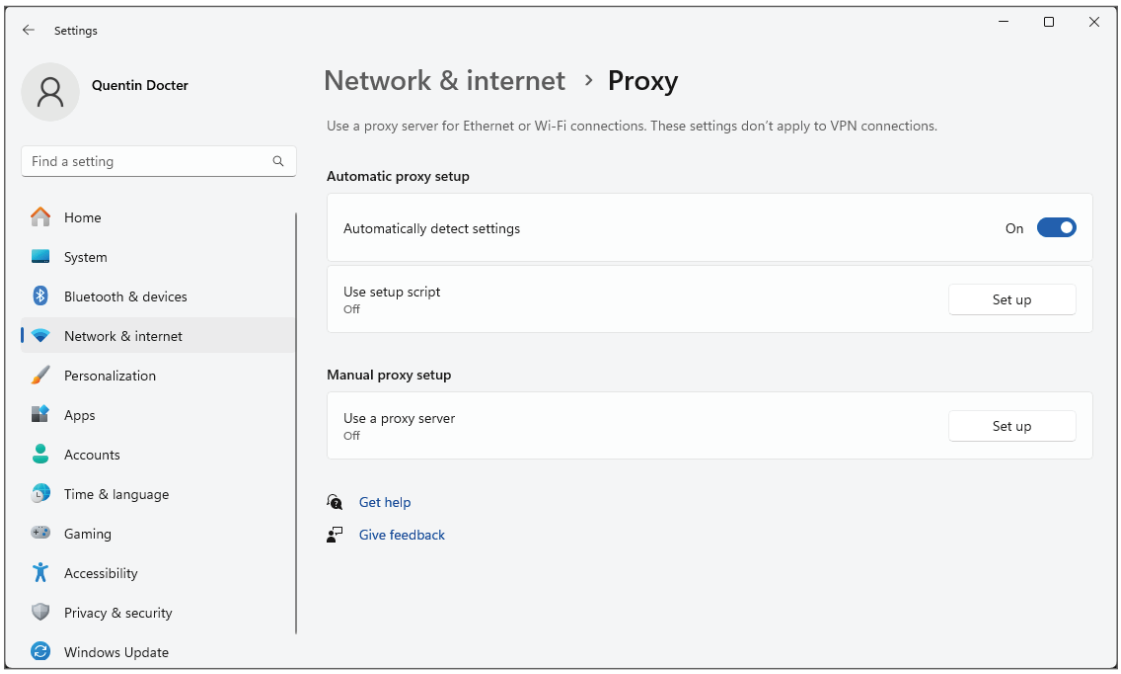

In the Automatic Proxy Setup section, there's an option to use a setup script. These can be used to reconfigure client browsers without manual intervention. Most often a setup script will provide configuration data so the client can properly use the proxy server.

4. To manually configure a proxy server, click the Set Up button next to Use A Proxy Server.

5. In the Edit Proxy Server window (Figure 8.17), enter the proxy server's IP address and the port it's configured to use. Notice you can also configure the proxy to exclude addresses from its services, and there is an option to also keep it from proxying intranet addresses.

EXERCISE 8.1 *(continued)*

FIGURE 8.17 Edit Proxy Server

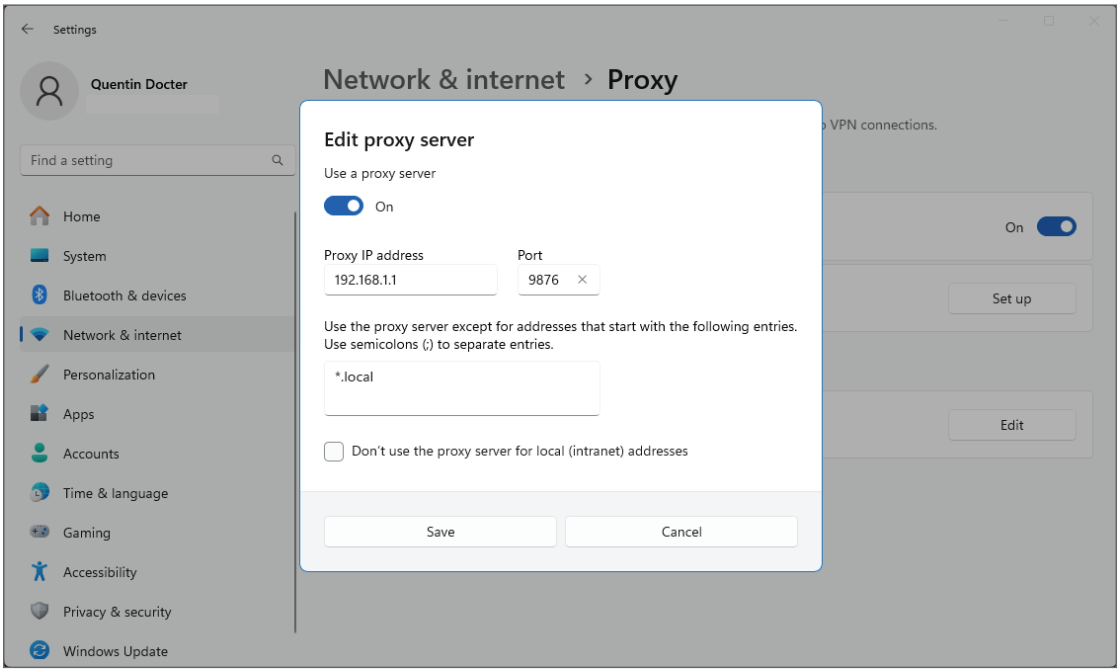

6. When you are done, click Save.

Let's use an example of a web proxy to illustrate how the proxy server process works. The user on the client computer opens a web browser and types in a URL. Instead of the request going directly to that website, it goes to the proxy server. The proxy then makes the request of the website, and it returns the requested information to the client computer. If it sounds to you like this slows down Internet browsing, you're right—it does. But there are three strong potential benefits to using a proxy.

Exam Objective 2.3 lists proxy servers as a subobjective. Understand the service they provide to network users.

First, the proxy server can cache the information requested, speeding up subsequent searches. (This is also the only function of a caching proxy, but caching-only proxies are most commonly configured to work on a local intranet.) Second, the proxy can act as a filter, blocking content from prohibited websites. Third, the proxy server can modify the

requester's information when passing it to the destination, blocking the sender's identity and acting as a measure of security; the user can be made anonymous.

Keep in mind that if all of the traffic from a network must pass through a proxy server to get to the Internet, that can really slow down the response time. Make sure the proxy or proxies have ample resources to handle all the requests. Figure 8.18 shows an example of a proxy server on a network.

FIGURE 8.18 A proxy server on a network

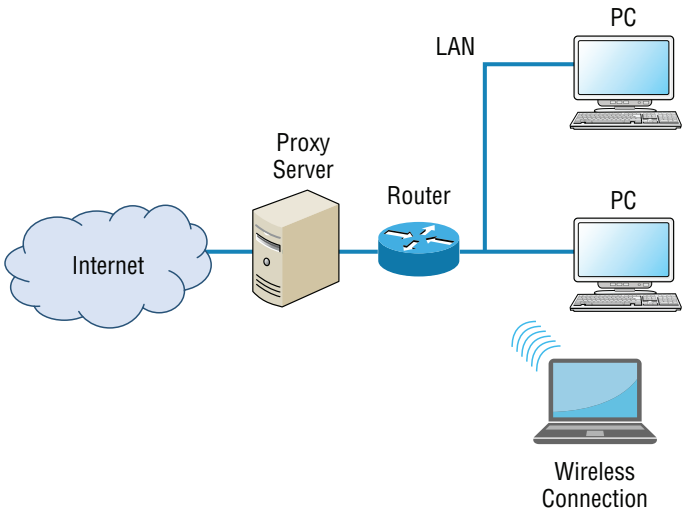

Legacy/Embedded Systems

In regular human terms, legacies are considered a good thing. Most of us want to leave a legacy of some kind, whether it's within a community or organization or within our own families. A legacy is something that lives on far beyond a human life span.

Exam Objective 2.3 lists legacy/embedded systems as a sub-objective. This includes a specific call out for SCADA. Know what these terms mean and how they affect network administration.

If you mention the term *legacy system* in the computer world, though, you are likely to be met with groans and eye rolls. It means that the system is old and hopelessly outdated by today's computing standards. Legacy systems are usually defined as those using old technology in one or more of the following areas:

- Hardware
- Software (applications or operating system; *legacy software*)
- Network protocols

Many legacy systems were state of the art when they were originally implemented in the 1970s or 1980s, but they haven't been upgraded or replaced. Today, though, they are old and slow, and specialized knowledge is required to maintain and operate them. For example, someone might need to know the Pick operating system (which came out in the 1970s), how to operate an IBM AS/400 or manage VAX, or how to configure the IPX/SPX network protocol. (Google these topics sometime!)

It's not just the really old stuff, though. Even technologies that have been invented after the turn of the century can now be considered legacy. For example, Microsoft no longer supports the Windows XP (released in 2001) and Windows 7 (2009) operating systems. Windows 10, which launched in 2015, will no longer be supported as of October 2025. A lot of computers still use Windows 10! Security and other updates will no longer be provided for these OSs, which could introduce security risks. On the hardware side, the wireless networking 802.11b standard is definitely legacy today, and if 802.11g isn't already legacy, it should be. It's really hard to find new components that include these standards.

An *embedded system* is one that is critical in a process; other systems or processes depend on it. For purposes of this section, we will treat legacy and embedded systems as one and the same, because administrators face similar issues with both.

So why don't companies replace legacy systems? It's complicated:

- First, most companies don't have large IT budgets, and replacing legacy systems can be very expensive. This is especially true for many companies that see the systems providing reliable (if a bit slow) service. Why fix what's not broken?
- Second, the cost of failure of an upgrade could be catastrophic. The world's global financial systems are in many places supported by legacy systems. Messing up a migration in that context could be a career-limiting move.
- Third, the time it would take to test the new system, verify functionality, and roll out the implementation could be extensive. Time is money, and we already stated that IT budgets are generally tight.

Furthermore, it's challenging to find technicians and consultants who understand legacy systems. People move from company to company, or consultants retire and take their specialized knowledge with them. Someone who was a mid-20s computer whiz in 1975 is now in their 70s and probably retired. The cost to find someone knowledgeable on these systems can be high.

Speaking of high costs, finding replacement hardware can be difficult to impossible while being expensive at the same time. Eventually, the cost of maintenance might outweigh the cost of upgrading, but then again it might not.

A great example of critical legacy systems is a category known as *supervisory control and data acquisition (SCADA)*. SCADA systems are high-level management systems that are used to control manufacturing machines and processes; manage large-scale infrastructure settings, such as power grids, oil and gas pipelines, and water treatment facilities; and

run components in buildings, such as heating and air conditioning. In other words, they're everywhere, and they manage some very important things. But, most SCADA systems are extremely old and were designed to be open access, so they are huge security holes. Hackers have been exploiting those holes faster than developers have been able to patch them.

So what's a network administrator to do? If possible, replacing or repurposing legacy systems can provide long-term benefits to a company. But also recognize the risk involved. If replacement isn't an option, then the best advice that we can give is to learn as much as you can about them. Hopefully, the system is based on established standards, so you can look them up on the Internet and learn as much as possible. If not, see what operating manuals you can track down, or pick the brain of those who understand how they operate. As challenging as legacy systems can be, you can make yourself quite valuable by being the expert on them.

A common administrative option is to try to isolate the legacy system as much as possible so that its lack of speed doesn't affect the rest of the network. It also increases security, because SCADA and industrial control systems (ICSs) can be more easily hidden within the network and newer systems will be more isolated from attempted attacks on them. This is usually much easier to do with hardware or protocols than with software. For example, the network might be set up with one segment that hosts the legacy systems or protocols.

A common administrative option is to try to isolate the legacy system as much as possible so that its lack of speed doesn't affect the rest of the network. This is usually much easier to do with hardware or protocols than with software. For example, the network might be set up with one segment that hosts the legacy systems or protocols.

One technology that is helping replace and update legacy systems is virtualization, which can obviate the need for one-to-one hardware-to-software relationships. We cover virtualization in Volume 1 Chapter 11.

Internet of Things Devices

The *Internet of Things (IoT)* has been growing rapidly for the last several years. Home automation systems and IoT-based security systems are becoming increasingly popular and will likely continue to become more pervasive. IoT also has a place in the manufacturing world, as companies seek to contain costs.

IoT networks will often have a central controller or coordinating device, like a computer switch but dedicated specifically to IoT devices. The administrator will have an app, usually on a smartphone, that provides them with a Wi-Fi or Bluetooth connection to the controller. Settings the administrator enters are then communicated to the end devices via the controller. Many IoT devices can be configured manually as well, but what's the fun in that? One of the key features of IoT is to control devices and see their status without needing to be physically present at each one. It could take an entire book to cover the different types of IoT devices, but we'll look at few popular ones here.

Thermostats

A thermostat is a device connected to a heating or cooling system that allows users to control the temperature of a room, home, or building. Programmable thermostats that allow you to set the temperature based on time and day have been around for more than 20 years. The next step in the evolution is a smart thermostat (shown in Figure 8.19) that's remotely accessible and can do some basic learning based on the weather and your preferences. Smart thermostats usually have a touch screen, often have their own app, and can be controlled by a central coordinator. Popular brands include Google Nest, Amazon Smart, ecobee, LG, and Honeywell.

FIGURE 8.19 ecobee smart thermostat

source: ecobee

Smart thermostats have the options you would expect, such as setting the temperature, detecting humidity levels, and configuring schedules to save energy when no one is home. The advanced features are what make these devices especially interesting. Examples include the following:

- Remote sensors, which can be placed on air vents for more accurate reading and control
- The capability to create energy reports
- Voice activation
- Integration with outdoor weather conditions to adjust internal conditions appropriately
- Geofencing, which allows the thermostat to adjust the temperature based on who it detects in the home and their preferences

Different models offer different features, so you're sure to find one that meets your needs as well as your budget.

Home Security and Automation

Home security systems have been around for a long time, and security cameras are an integral part of them. With IoT-based cameras, security systems are more customizable and smarter than ever.

The security camera can be stand-alone, but most often it is part of a series of cameras connected to a home security system. That central system will often have a touchscreen as well as an app. The system can be activated in several ways, such as by the ringing of an integrated doorbell or by motion sensors. Most systems will provide for the recording of video in the event that they're triggered by motion. Video footage may be stored on local storage, such as an SD card, or on the cloud. Configuration options often include when to activate (e.g., when motion is detected), how to notify the user (such as a text or email), and how long to record and store footage.

Important features to consider when buying a smart security camera include its movement capability, field of view, resolution, night vision, waterproofing, audio capabilities (microphone and/or speakers), and motion detection. Cost for the monitoring service and how much storage capacity the service offers may also be a factor.

Home security and automation systems may also control door locks and light switches.

Door Locks

Smart door locks are a feature often integrated into home security systems. They will typically be accompanied by a camera and linked to a doorbell. Many will also have a number pad so that you can enter a code to unlock the door instead of needing a key. Figure 8.20 shows a Schlage smart door lock.

FIGURE 8.20 Schlage smart door lock

In addition to offering security, smart door locks can provide convenience. For example, perhaps you are expecting a delivery. Someone rings the doorbell. With a smart security system that includes a doorbell and door locks, you can instantly see who it is on your smartphone and talk to them. If you feel comfortable, you can unlock the door and tell them to place the package inside. Once they are done, you can lock the door again. Or, you can tell them to set the package down and you'll come get it in a few minutes, whether you're on the other side of the house or the other side of the country.

Light Switches

Smart light switches help control lights in the house. Many are designed to replace existing light switches in the wall, whereas others simply mount to the wall. An example of a Lutron switch is shown in Figure 8.21. In addition to having manual controls, many will have their own app or can be controlled through a coordinator.

FIGURE 8.21 Lutron smart light switch

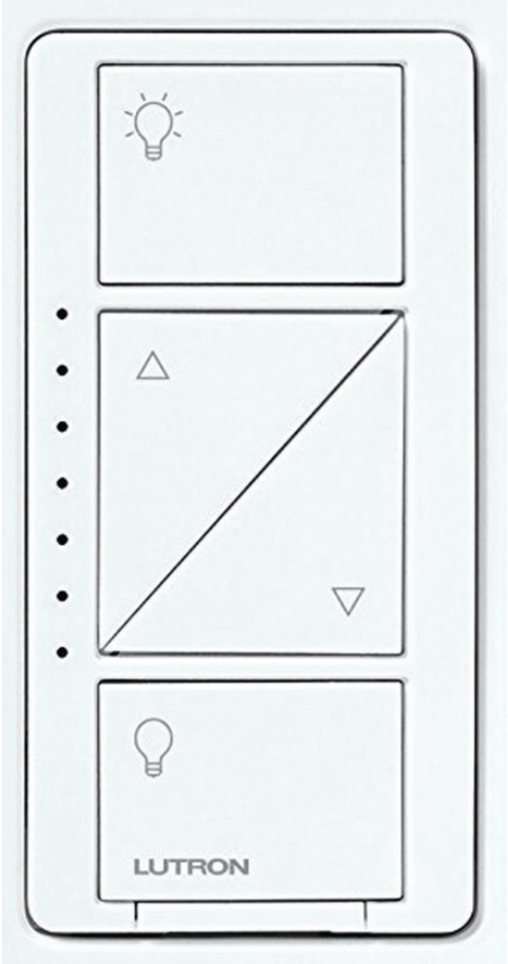

Features of a smart light switch are fairly straightforward. They can turn lights on or off and dim the lights. They can perform tasks based on a schedule, and some have geofencing or motion sensors to detect when someone enters a room. Some brands will work only with certain types of lights, so make sure to check compatibility.

Voice-Enabled Smart Speakers/Digital Assistants

Smartphones ushered in the widespread use of voice-enabled digital assistants. It started with Siri on the iPhone, and Google Now ("Okay, Google") soon followed for the Android OS. Microsoft even got into the act with Cortana, which was used with its now defunct Windows Phone OS and was also integrated into Windows 10. (It's been depreciated in Windows 11.) Amazon wanted in on the act, too, but they don't have a smartphone OS. So instead, they created a voice-enabled smart speaker called the Echo with a virtual assistant known simply as Alexa. Google, not to be outdone in the digital assistant market, created Google Home, which uses Google Assistant, which evolved from Google Now. The market for these devices is very competitive.

Siri is proprietary to Apple devices, but smart speakers made by other companies will support Alexa or Google Assistant, or both.

The first smart speakers on the market were Wi-Fi–enabled speakers that would listen for you to activate them. Once you said their name, they would listen to your question and use their Internet connection to perform a search and deliver an answer. They were incredibly handy for asking the weather for the day, to play a song, or to answer an obscure trivia question at a party. Newer smart speakers may have integrated video screens as well. Most are in the 8″ or 10″ range, but some are up to 21″ in size. This makes them larger and more conspicuous but lets users view their integrated security cameras, see music videos and movies, and more.

While the features of a smart speaker/digital assistant are appealing to many, there are some potential concerns and risks as well. For example, unless you turn it off, the device is always listening. This is unnerving to some people, and others have even suggested it could be used to eavesdrop on its users. Device manufacturers have gone to lengths to protect users' privacy but know that this could be an issue.

A famous story of misuse comes from late-night television. During a late-night show in 2017, comedian Jimmy Kimmel decided to pull a prank on all of his viewers who had an Amazon Echo or similar device. With the audience quiet, he loudly instructed Alexa (the name that the Amazon digital assistant responds to) to order $500 worth of foam swim "pool noodles." Urban legend has it that multiple viewers were affected by this and sent the product. (To be fair, there are conflicting reports of whether or not it actually worked. If it did, Amazon would have accepted the return and reversed the charge.) Of course, this problem could be avoided by turning off automatic voice ordering from Amazon, which is a feature of its digital assistants. Finally, some flaky smart speaker units have been known to speak when not spoken to or even to start laughing for no apparent reason. While this is definitely creepy, it's not necessarily a security threat.

Summary

In this chapter, you learned about different server roles and technologies that work to provide services to networked clients.

First, you learned about specific server roles. The two we went into the most depth on are DNS and DHCP. DNS resolves hostnames to IP addresses through the use of a zone file. It's configured to query other DNS servers if it isn't able to resolve the address on its own. We talked about record types including A, AAAA, CNAME, MX, and TXT. TXT is especially important in reducing spam emails through DKIM, SPF, and DMARC.

DHCP provides TCP/IP configuration information to client computers. It uses a scope to define the addresses available, and clients get the IP address on a timed lease. Scopes can also have reservations and exclusions.

Other server types/roles we covered include file (fileshare), print, mail, syslog, web, AAA, database, and NTP servers. We talked about what each one of these does as well as where they should be located on the network, either inside the secure network or in the screened subnet. In addition to servers, many networks will have Internet appliances dedicated to security, such as spam gateways and UTM devices, as well as load balancers and proxy servers to manage traffic. Some networks also support legacy or embedded systems, such as SCADA. Although these systems are old and outdated, they often provide critical functionality on the network.

Finally, we looked at some examples of IoT devices. IoT devices are commonly used as thermostats and as part of home security systems. Voice-enabled assistants such as Amazon Alexa, Apple's Siri, and Google Assistant provide a hands-free experience.

Exam Essentials

Know the various roles that servers can play on a network. Roles include DNS, DHCP, file (fileshare), print, mail, syslog, web, AAA, database, and NTP servers. File servers (fileshares) store files for users, and may have optical media and perform backups too. Print servers host printers. Mail servers store, send, and receive email. A syslog server is used to log system events. Web servers host web pages that users access across a network. AAA servers validate user credentials, and then allow users to access resources and track access. Database servers host databases and use a DBMS. NTP servers synchronize time across a network.

Know what DNS servers do. DNS servers resolve hostnames to IP addresses. Without DNS servers, finding your favorite websites on the Internet would be an incredibly challenging task. DNS servers have a zone file with hostname to IP address mappings.

Know five common DNS record types and what each one does. DNS record types include A, AAAA, CNAME, MX, and TXT. A is for IPv4 hosts and AAAA (quad A) is for IPv6 hosts. CNAME is an alias, MX is a mail server, and TXT is a text entry.

Understand how DNS helps mitigate spam email. DNS uses DKIM, SPF, and DMARC to help fight email spam.

Understand how DHCP servers work. DHCP servers lease IP addresses and configuration information to client computers. Clients request the information via broadcast. Each DHCP server has a scope with a configured range of available IP addresses. The server may also provide additional configuration information, such as the address of the default gateway (a router) and a DNS server. Scopes may also have exclusions and reservations.

Understand what spam gateways and UTM systems do. Spam gateways help email servers detect and deal with unwanted spam email. Unified threat management (UTM) systems centralize security management and often replace traditional firewalls.

Know what load balancers and proxy servers do. Both types of servers help manage network traffic. Load balancers do so by sending incoming requests to different, typically identical servers to spread out the workload. Proxy servers make requests on behalf of clients.

Know what legacy and embedded systems are. Legacy systems are older technology no longer supported by the manufacturer. Embedded systems are those that are critical to a process. SCADA is an example of a legacy and embedded system.

Know some examples of services provided by Internet of Things (IoT) devices. Some device types include thermostats, home security and automation, and voice-enabled speakers and digital assistants.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. Your company hosts its own web server, and it allows consumers to make purchases via the server. The help line has been getting complaints that users are unable to access the website. You open the site from an internal workstation and it seems fine. What is the most likely cause?
   - **A.** The firewall is blocking TCP port 23.
   - **B.** The firewall is blocking TCP port 443.
   - **C.** The security module of the web server is malfunctioning.
   - **D.** The web server is down.
2. Your manager wants you to install a networked Internet appliance that prevents network traffic–based attacks and includes antimalware and antispam software. What should you install?
   - **A.** Spam gateway
   - **B.** Load balancer
   - **C.** UTM
   - **D.** Proxy server
3. You are installing a file server for the accounting department. Where should this file server be located on the network?
   - **A.** Outside of the firewall
   - **B.** In the screened subnet
   - **C.** In the secure network
   - **D.** On the router
4. Which of the following are services that a print server should provide? (Choose two.)
   - **A.** Accepting print jobs from clients
   - **B.** Turning off printers on demand
   - **C.** Providing clients with the appropriate printer driver during installation
   - **D.** Notifying users when the print job is complete
5. Which type of server is responsible for preventing users from accessing websites with objectionable content?
   - **A.** Proxy
   - **B.** Web
   - **C.** DHCP
   - **D.** DNS

6. What does a DHCP server need to be configured with to operate properly?
   - **A.** DNS server
   - **B.** Scope
   - **C.** Range
   - **D.** DHCP relay agent

7. When configuring a DNS server, administrators must create which of the following?
   - **A.** Zone file
   - **B.** Hosts file
   - **C.** Scope file
   - **D.** DNS proxy

8. You are configuring two email servers on your company's network. Which network protocol do the servers use to transfer mail to each other?
   - **A.** POP3
   - **B.** IMAP4
   - **C.** SNMP
   - **D.** SMTP

9. You have five web servers that manage requests for online purchases. An administrator notices that one of the severs is always busy while another is idle, and the company is getting some online complaints about the slow website. Which of the following servers will help fix this?
   - **A.** DNS
   - **B.** DHCP
   - **C.** Proxy
   - **D.** Load balancer

10. A computer using which of the following would be considered a legacy device? (Choose all that apply.)
    - **A.** A 386 processor
    - **B.** The IPX/SPX protocol
    - **C.** An application developed in 1983
    - **D.** 1 GB of RAM

11. You are configuring a DHCP server for a small network with a network address of 192.168.1.0/24. You want the DHCP server to manage and assign all IP addresses and also want to ensure that there are 20 addresses available for servers, printers, and router interfaces. Which of the following is the best solution?
    - **A.** Create a scope with addresses 192.168.1.1 to 192.168.1.200. Create an exclusion for addresses 192.168.1.1 to 192.168.1.20.
    - **B.** Create a scope with addresses 192.168.1.1 to 192.168.1.200. Create a lease for addresses 192.168.1.1 to 192.168.1.20.

C. Create a scope with addresses 192.168.1.21 to 192.168.1.200. Create an inclusion for addresses 192.168.1.1 to 192.168.1.20.

D. Create a scope with addresses 192.168.1.1 to 192.168.1.200. Create reservations for addresses 192.168.1.1 to 192.168.1.20.

12. You're adding entries for new servers into the corporate DNS server and have IPv6 addresses for those servers. Which type of entries will you create?

A. A

B. CNAME

C. AAAA

D. TXT

13. An administrator sets up an NTP server on their local network. What is the purpose of an NTP server on the local network?

A. To query national NTP servers to ensure all computers on the local network are set to the same time

B. To provide the correct time to all computers on the local network

C. To ensure all local computers have the same newsfeed

D. To resolve SCADA control issues between compliant devices

14. Company employees have recently started receiving large amounts of spam emails. Which of the following systems can help reduce the amount of spam received? (Choose all that apply.)

A. UTM

B. DMARC

C. SCADA

D. AAA

15. A network administrator is updating features on an AAA server. Which aspect of that server deals specifically with determining which users get to access which network resources?

A. Authentication

B. Authorization

C. Accounting

D. All of the above

16. Which type of network service is responsible for collecting and tracking error messages generated by networked devices such as routers?

A. AAA

B. NTP

C. Syslog

D. UTM

17. Which of the following is a term that refers to legacy industrial servers that control critical infrastructure?
    - **A.** DKIM
    - **B.** SPF
    - **C.** DMARC
    - **D.** SCADA

18. Your company has decided to set up its own DHCP server. How many DHCP servers is the company required to maintain?
    - **A.** 1.
    - **B.** 2.
    - **C.** 3.
    - **D.** It depends on how many computers are on the network.

19. You just purchased an Alexa smart assistant. What type of service is this?
    - **A.** UTM
    - **B.** SCADA
    - **C.** IoT
    - **D.** AAA

20. Which spam management capability authenticates email servers based on their IP addresses?
    - **A.** SPF
    - **B.** DKIM
    - **C.** DMARC
    - **D.** SCADA

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answer compares to the authors', refer to Appendix B.

You need to configure a network for the proper amount of security that still meets all user needs, and are given the hardware shown in the image. Draw what the network should look like, if using a two-pronged firewall approach, and locate the servers and clients appropriately.

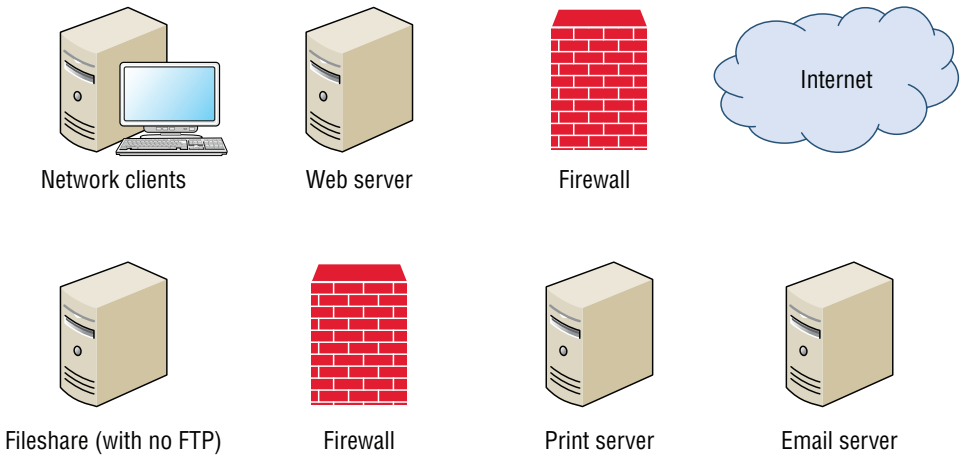

Chapter 9

Wireless and SOHO Networking

THE FOLLOWING COMPTIA A+ 220-1201 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

✓ **2.0 Networking**

✓ **2.2 Explain wireless networking technologies.**

- Frequencies
  - 2.4GHz
  - 5GHz
  - 6GHz
- Channels
  - Regulations
  - Channel selection
    - Widths
    - Frequencies
    - Bands
- Bluetooth
- 802.11 standards
- NFC
- Radio-frequency identification (RFID)

✓ **2.7 Compare and contrast internet connection types, network types, and their characteristics.**

- Internet connection types
  - Satellite
  - Fiber

- Cable
- DSL
- Cellular
- Wireless internet service provider (WISP)

Over the last three chapters, we've talked a lot about foundational networking knowledge. We've discussed theoretical networking models, physical topologies, cables and connectors, and connectivity devices. We also spent an entire chapter devoted to the most common protocol of all, TCP/IP. One critical technology that we haven't covered yet is wireless networking.

Because of the unique technology of wireless networking and its huge popularity, it feels appropriate to talk about it as a separate entity. That said, it's important to remember that wireless networking is just like wired networking, only without the wires. You still need to figure out how to get resources connected to each other and give the right people access while keeping the bad people at bay. You're now just playing the game with slightly different rules and many new challenges.

This chapter serves two purposes. The first is to finish out the last of our networking "theory" discussions, this time on the categories of wireless networking standards. The second is to put all of the networking knowledge you've gained so far into one practical picture, which is how to set up a small office/home office (SOHO) network. We'll start by covering common Internet connection types, because all networks today need an Internet connection. After that, we'll walk through important concepts you need to consider when setting up a network. Particular attention will be given to setting up a wireless router, which is the most prevalent networking tool found on SOHO networks today. At the end of this chapter, you should have full confidence in being able to set up a SOHO network in real life, as well as being able to pass the networking portion of the A+ exam.

Understanding Wireless Networking Technologies

Wireless networking is so common today that it's taken for granted. When first introduced, wireless networking was slow and unreliable, but it's now fast and pretty stable, not to mention convenient. It seems Internet cafes are everywhere and nearly all restaurants have wireless hotspots. Every mobile device today seemingly has Internet capabilities—if it doesn't, nobody will buy it and it will disappear quickly. No matter where you go, you're likely just seconds away from being connected to the Internet.

The most common term you'll hear thrown around referring to wireless networking today is *Wi-Fi*. While the term was originally coined as a marketing name for 802.11b, it's now used as a nickname referring to the family of IEEE 802.11 standards. That family comprises the primary wireless networking technology in use today, but other wireless technologies are out there, too. We'll break down wireless technologies into three groups: 802.11, Bluetooth, and radio frequency. Each technology has its strengths and weaknesses and fills a computing role.

As a technician, it will fall to you to provide users with access to networks, the Internet, and other wireless resources. You must make sure that their computers and mobile devices can connect, that users can get their email, and that downtime is something that resides only in history books. To be able to make that a reality, you must understand as much as you can about wireless networking and the technologies discussed in the following sections.

Understanding 802.11 Networking

In the United States, wireless LAN (WLAN) standards are created and managed by the Institute of Electrical and Electronics Engineers (IEEE). The most commonly used WLAN standards are in the IEEE 802.11 family. Eventually, 802.11 will likely be made obsolete by newer standards, but that is some time off. IEEE 802.11 was ratified in 1997 and was the first standardized WLAN implementation. There are dozens of 802.11 standards defined, but you will likely only hear a few commonly mentioned: 802.11ac, ax, and be. As previously mentioned, several wireless technologies are on the market, but 802.11 is the one currently best suited for WLANs.

Along with the IEEE, another important organization in the Wi-Fi space is the Wi-Fi Alliance (`www.wi-fi.org`). IEEE develops and ratifies the 802.11 standards. Then the Wi-Fi Alliance takes the standards and develops "real world" principles for applying them. This includes testing and certifying devices to confirm that they operate within the standards, to ensure interoperability between device types and manufacturers. When looking for Wi-Fi devices, be sure that they are Wi-Fi certified.

In concept, an 802.11 network is similar to an Ethernet network, only wireless. At the center an Ethernet network is a connectivity device, usually a switch, and all computers are connected to it. Wireless networks are configured in a similar fashion, except that they use a wireless router or wireless access point instead of a wired connectivity device. In order to connect to the wireless hub or router, the client needs to know the *service-set identifier (SSID)* of the network. SSID is a fancy term for the wireless network's name. Wireless access points may connect to other wireless access points, but eventually they connect back to a wired connection with the rest of the network.

802.11 networks use the *Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA)* access method instead of Ethernet's Carrier Sense Multiple Access/Collision Detection (CSMA/CD). Packet collisions are generally avoided, but when they do happen, the sender will need to wait a random period of time (called a *back-off time*) before transmitting again.

Since the original 802.11 standard's publication in 1997, several upgrades and extensions have been released. The primary characteristics that define them are speed and maximum distance.

Defining Wi-Fi Terms

Before we get into specific Wi-Fi standards and their differences, its important to set the foundation with key terms you need to understand. All of these terms will be used later in relation to how the standards integrate or upgrade them, resulting in better performance from generation to generation.

Frequency All wireless networking technologies, by definition, send signals through the air. The *frequency* is the range of radio waves used to transmit data. It's measured in gigahertz (GHz), which is how many billions of times a full wave oscillates through a given point in a second. The three frequencies that Wi-Fi uses (depending on the version) are 2.4 GHz, 5 GHz, and 6 GHz. You will also hear these referred to as *bands*, for example, the 2.4 GHz band versus the 5 GHz band.

Lower frequencies can't carry as much data as higher ones, but they do have longer range and are less susceptible to interference from obstructions such as walls. 6 GHz signals, for example, tend to only perform best with line-of-sight between two devices. This isn't to say that 6 GHz signals *can't* go through walls or other obstructions—they can. It's just that they get more severely degraded than lower frequency waves.

Exam Objective 2.2, "Explain wireless networking technologies," calls out frequencies and channels as subobjectives. For example, you should know general performance characteristics about 2.4 GHz, 5 GHz, and 6 GHz frequencies. Specifically, remember that lower frequencies carry less data, but have longer range and are less susceptible to interference. The "best" technology to use will often depend on the circumstances.

You should also know channel widths, frequencies, bands, and regulations around channels and channel selection. We're defining each of the terms here, but that knowledge alone won't be enough for the exam. We will also explain them in more detail as we talk about the different 802.11 standards, which will help you for the exam.

Channels Each Wi-Fi frequency is broken down into *channels*, or narrower bands of frequencies. If the frequency is analogous to a highway, think of channels as lanes on the highway. Older Wi-Fi implementations used single channels, whereas newer ones can combine channels to increase throughput.

Each channel has a *width*, which defines its band of frequencies. The number of channels is different per band. Table 9.1 outlines the basics of frequencies and their maximum channels and band width.

TABLE 9.1 Frequencies and channels

Frequency	Channel band width	Max 20 MHz channels
2.4 GHz	20 MHz	14
5 GHz	20 MHz, 40 MHz, 80 MHz, or 160 MHz	40-45
6 GHz	20 MHz, 40 MHz, 80 MHz, or 160 MHz	59

The maximum number of 5 GHz channels depends on the region and regulations within that region. Weather, commercial, and military radar may operate in this range as well, causing the variance. Therefore, there may be regulations on which channels can be used for a Wi-Fi network.

Modulation Modulation refers to how the Wi-Fi signals are modified for transmission between the sender and receiver to convey data. Typical types of modulation change the frequency (speed), amplitude (strength or height), or phase (position at a point in time) of the signal. Both the sender and the receiver need to use the same modulation type to communicate. Newer modulation methods allow for higher data transfer speeds. In the "Influencing the Signal" sidebar later in this chapter, we get into more technical details.

Coding Scheme Wi-Fi has built-in error detection, and it's implemented in the form of a coding scheme. Specifically, the coding scheme defines how many data bits and error correction bits are transferred together.

Spatial Streams Newer Wi-Fi devices have multiple antennas and can send and receive data simultaneously on these different antennas. A spatial stream is a single data transmission that's sent using multiple antennas. Combining the antennas into one stream allows for increased throughput.

Physical Rate The physical rate, or PHY rate, is the theoretical maximum data throughput of a Wi-Fi standard. Notice we said theoretical, because PHY rates are never achieved in the real world. PHY rates are determined by factors such as frequency, modulation, channel width, coding scheme, and spatial streams.

Data Rate The data rate is the real throughput of a Wi-Fi connection. In addition to the characteristics used to calculate the PHY rate, data rate will depend on factors such as distance between devices, network overhead (how many devices are on the network), guard interval (how long the devices pause before sending more data) and interference. For example, it's no secret that the further away from an access point you are, the slower the Wi-Fi signal will be. The same is true if one thousand people are trying to access a Wi-Fi network versus just one person. We've all experienced these conditions.

Going deeper into the weeds, the maximum data rate is continually calculated by sending and receiving Wi-Fi devices in order to maximize data throughput. Wi-Fi providers and engineers use the Modulation Coding Scheme (MCS) index to determine the

maximum data rate between two devices. You're unlikely to be tested on the MCS, and its calculations are relatively complex. If you're curious and want to see a comprehensive MCS index table outlining maximum speeds based on connection characteristics, go to `https://mcsindex.net`.

Now that we have those terms under our belts, let's dive into 802.11 standards.

Legacy 802.11 Standards

The original *802.11* standard was ratified in 1997 and defines WLANs transmitting at 1 Mbps or 2 Mbps bandwidths using the 2.4 GHz frequency spectrum. The frequency-hopping spread spectrum (FHSS) and direct-sequence spread spectrum (DSSS) modulation techniques for data encoding were included in this standard.

There were never any 802.11 (with no letter after the 11) devices released—it's a standards framework. All in-market versions of 802.11 have a letter after their name to designate the technology. We'll introduce these to you so we can show you how far technology has come in a relatively short time and introduce key Wi-Fi concepts and features that are relevant for today's versions as well.

Exam Objective 2.2 tells you to be familiar with 802.11 standards, but it doesn't specify which ones. You should have some basic familiarity with these legacy standards, just in case.

802.11a

The 802.11a standard provided WLAN bandwidth of up to 54 Mbps in the 5 GHz frequency spectrum. The 802.11a standard also used a more efficient encoding system, orthogonal frequency division multiplexing (OFDM), rather than FHSS or DSSS.

This standard was ratified in 1999, but devices didn't hit the market until 2001. Thanks to OFDM, it was significantly faster than 802.11b (discussed next) but never gained widespread popularity. They were ratified as standards right around the same time, but 802.11b devices beat it to market and were significantly cheaper.

Influencing the Signal

If you're someone who likes getting into details about how electromagnetics work, modulation techniques are a cool area to study. If you're not, that's okay because the A+ exam won't likely test you on them. Knowing that you're unlikely to get questions on the exam about specific ones, we're going to dive in anyway to give you a quick primer. There's a method to our madness—you might not get a question on OFDM versus FHSS, but understanding how and why they are different might help you answer a tough networking question, and it can also help you in the real world to understand why two Wi-Fi devices might refuse to talk to each other.

So far we have introduced three modulation techniques: FHSS, DSSS, and OFDM. All of these are techniques used to get data transmitted between two devices on a Wi-Fi network. Specifically, it determines how the sending device gets the signal to a receiving device. (Modulation is used in other wireless communications such as cellular and Bluetooth as well.) They differ in how they accomplish their goals, however:

- FHSS "frequency hops," or switches rapidly between frequencies.
- DSSS takes the signal and spreads it out across the available band.
- OFDM divides the signal into pieces and transmits it among multiple subcarriers.

There's also a fourth that's used called Orthogonal Frequency Division Multiple Access (OFDMA) that we haven't covered yet—it was introduced with Wi-Fi 6, which we haven't gotten to yet—that's like OFDM but across multiple antennas.

To complicate matters, there are modulation schemes that work within these techniques that define how the bits get encoded within the data stream (or streams). In other words, the modulation techniques we shared earlier are how the signal gets from point A to point B, whereas the schemes here are how the signal gets turned into data that makes sense. The three common ones in Wi-Fi are as follows:

- Binary Phase Shift Keying (BPSK) transmits one bit per symbol. Think of a symbol as representing data, either one bit or a group of bits, that's transmitted over a specific amount of time. Each symbol is marked by a unique combination of the wireless signal's phase and amplitude. BPSK is the simplest and uses the least amount of power among the three modulation schemes.
- Quadrature Phase Shift Keying (QPSK) transmits 2 bits per symbol by modifying the signal's phase.
- Quadrature Amplitude Modification (QAM) transmits multiple bits per symbol, giving it the highest data throughput. It's also the most sensitive to interference.

The sending device uses a technique (such as OFDM) that uses a scheme (such as QPSK) to convey the message to the receiving device.

This can seem pretty complicated—it might feel foreign to anyone without an electrical engineering degree. But remember that there's a method to our madness. There are three methods, to be exact. First, it's so you can better understand why two Wi-Fi devices need to use the same modulation techniques to communicate. If they don't, the receiving device has no idea what to do with the incoming message. Second, when you see the specifications for Wi-Fi standards, you might run into these terms. Third, some devices in the market use the terms too. For example, you might see a wireless router advertisement trying to impress you by saying it supports 4K-QAM. Now you know what QAM means, we'll get into 4K-QAM specifically when we introduce Wi-Fi 7 later in this chapter.

802.11b

The 802.11b standard was ratified in 1999 as was 802.11a, but device makers were much quicker to market, making this the de facto wireless networking standard for several years. 802.11b provided for bandwidths of up to 11 Mbps (with fallback rates of 5.5, 2, and 1 Mbps) in the 2.4 GHz range.

Remember that for interoperability purposes, 802.11*x* devices across standards can communicate with each other only if they speak on the same frequency and support the same modulation techniques.

802.11g

Ratified in 2003, 802.11g standard provided for bandwidths of 54 Mbps in the 2.4 GHz frequency spectrum using OFDM or DSSS encoding. Because it operated in the same frequency and could use the same modulation as 802.11b, the two standards were compatible. Because of the backward compatibility and speed upgrades, 802.11g replaced 802.11b as the industry standard for several years.

Changing the Channel

One additional concept to know about when working with 2.4 GHz wireless networking is channels. Within the 2.4 GHz range, the Federal Communications Commission (FCC) has defined 14 different 22 MHz communication channels. This is analogous to a 14-lane audio wave highway, with each lane being 22 MHz wide. An illustration of this is shown in Figure 9.1. (We're aware that previously in Table 9.1 we said they were 20 MHz channels. Technically they *are* 22 MHz. However, most devices that use the 2.4 GHz band categorize them as 20 MHz, and most admins think about them that way too. The difference is immaterial, and you're not going to get a test question with 20 MHz and 22 MHz as different answer options.)

FIGURE 9.1 2.4 GHz communication channels

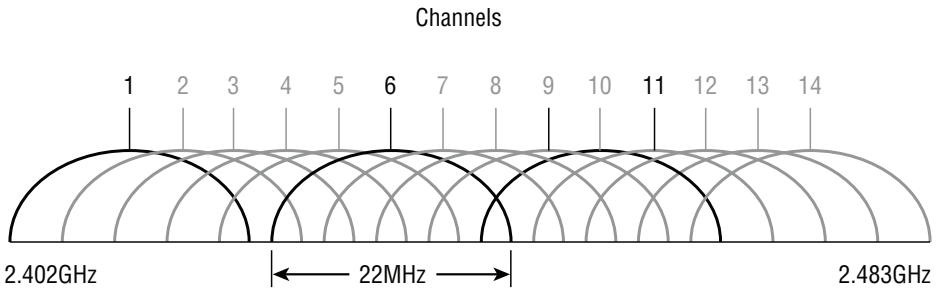

Although 14 channels have been defined for use in the United States, you're allowed to configure your wireless networking devices only to the first 11. When you install a wireless

access point and wireless NICs, they will all autoconfigure their channel, and this will probably work okay for you. If you are experiencing interference, changing the channel might help. This is done through the wireless router's configuration software, and then still have the NICs autodetect.

When setting up a Wi-Fi network with multiple access points, you should always have overlap between their ranges to ensure there are no dead spots in coverage. In situations like that, you need to have nonoverlapping channels to avoid communications problems. Figure 9.2 shows a few examples. Two channels will not overlap if there are four channels between them. If you need to use three nonoverlapping channels, your only choices are 1, 6, and 11. Notice back in Figure 9.1 that those three channels are in bold font.

FIGURE 9.2 2.4 GHz network layouts

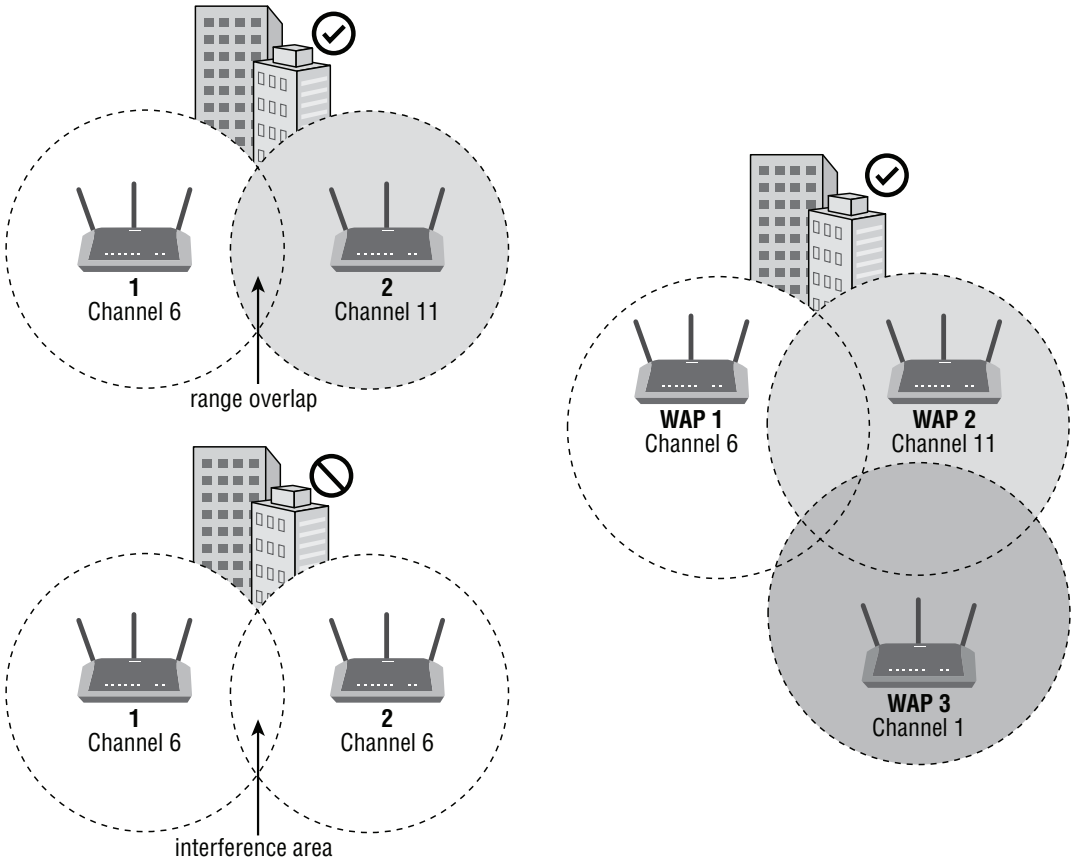

There are many other commercial devices that transmit in the 2.4 GHz range as well, such as Bluetooth devices, cordless phones, cell phones. Other WLANs and even microwave ovens can generate interference problems in this range. Changing the channel might help relieve interference with these devices.

802.11n

Next was 802.11n (retroactively named Wi-Fi 4), ratified in 2010. The standard had a PHY rate of 600 Mbps, but in reality the typical throughput was about 300–450 Mbps. It was still pretty fast, and it was backward compatible with 802.11b/g.

802.11n achieved faster throughput in a couple of ways. Some of the enhancements included the use of wider 40 MHz channels, multiple-input multiple-output (MIMO), and channel bonding. Remember how 802.11b/g used 20 MHz channels? 802.11n combined two channels to double (basically) the throughput. Imagine being able to take two garden hoses and combine them into one bigger hose. That's kind of what channel bonding does. MIMO means using multiple antennas rather than a single antenna to communicate information. (802.11n devices could support up to eight antennas, or four spatial streams, because each antenna only sent or received.) Channel bonding also allowed the device to communicate simultaneously at 2.4 GHz and 5 GHz and bond the data streams, which increased throughput. MIMO and channel bonding are employed in newer standards as well.

Not As Much Freedom As You Might Think

In the 5 GHz spectrum, there are 25 nonoverlapping 20 MHz communications channels, 24 of which can be used for Wi-Fi networks. On the surface, this sounds much better than the three nonoverlapping channels available in the 2.4 GHz spectrum.

To increase throughput, 802.11n bonded channels together. When 20 MHz channels are bonded into 40 MHz channels, this reduces the number of nonoverlapping channels to 12. To complicate matters further, weather, commercial, and military radar operate in the 5 GHz range as well, and Wi-Fi needs to avoid conflicting with them.

To avoid conflicts, wireless routers used a technology named *dynamic frequency selection (DFS)*, which detected radar interference and dynamically adjusted to a different frequency range to avoid the problem. If an installation was in an area with no interference from radar signals, it had 12 nonoverlapping 40 MHz channels; otherwise, only four nonoverlapping, non-DFS 40 MHz channels remained available for bonding.

If you're curious (and it's highly unlikely that you will be tested on this), the four nonoverlapping, non-DFS 40 MHz channels are numbered 36 and 40, 44 and 48, 149 and 153, and 157 and 161. We explain this a bit more in the upcoming section on 802.11ac, and Figure 9.3 illustrates what we just discussed.

FIGURE 9.3 Channel availability in the 5 GHz spectrum

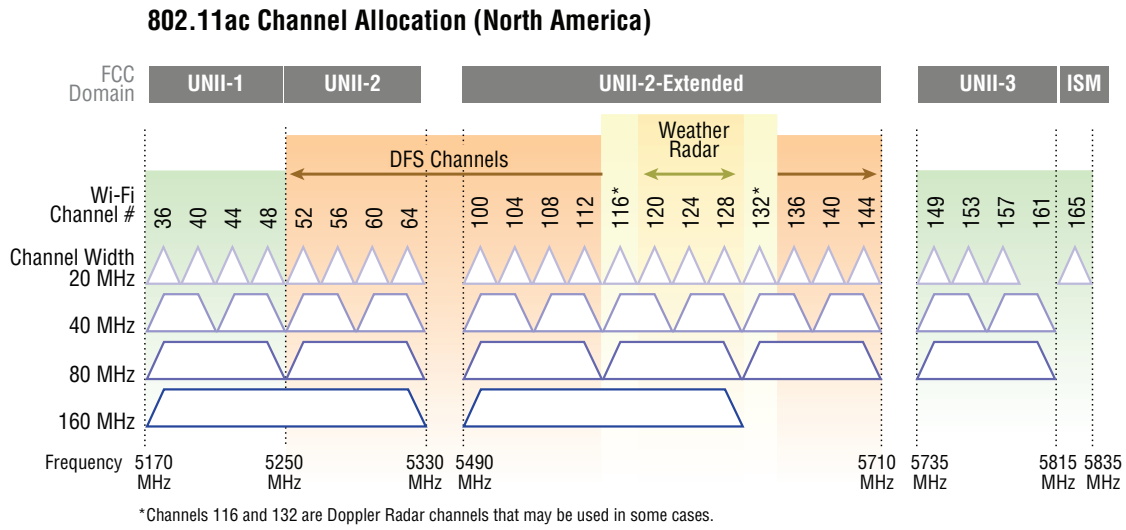

Table 9.2 provides a quick comparison of the legacy Wi-Fi standards we just covered.

TABLE 9.2 Legacy 802.11 standards

Type	Frequency	PHY rate (throughput)	Modulation	Indoor range	Outdoor range
—	2.4 GHz	2 Mbps	FHSS/DSSS	20 meters	100 meters
a	5 GHz	54 Mbps	OFDM	35 meters	120 meters
b	2.4 GHz	11 Mbps	DSSS	40 meters	140 meters
g	2.4 GHz	54 Mbps	DSSS/OFDM 64-QAM	40 meters	140 meters
n	5 GHz/2.4 GHz	600 Mbps	OFDM/DSSS 64-QAM	70 meters	250 meters

The Long and Short on Wi-Fi Ranges

The ranges provided in Table 9.2 are approximate and may differ based on your environment. For example, thick walls and steel beams will dramatically reduce your range. In addition, the maximum data rates for 802.11n and newer standards can be debated. Some equipment (and experts) will quote the theoretical maximum PHY rate, whereas others will

give more realistic throughput numbers. This wasn't an issue pre-802.11n, but the newer technologies use fancier modulation techniques to achieve maximum speeds, and we don't always have hardware that can produce the speeds that scientists can draw up on a whiteboard.

Also keep in mind that when discussing ranges, the farther away from the WAP you get, the lower your connection speed will be. For example, to get 54 Mbps out of an 802.11g router, you need to be within about 100 feet of it. At the far end of its range, your throughput will be only about 6 Mbps. Another key is how many clients are attached to the WAP. More clients mean less bandwidth for each client. If you have ever used the Wi-Fi in an airport or a busy hotel, you know exactly what we mean.

Gigabit 802.11 Standards

Technology is always marching forward and getting faster and cheaper, and wireless networking is no different. In previous chapters you've heard us say that when it comes to bus speeds (such as USB and PCIe), the next generation doubles the speed of the previous one. That's not necessarily the case with Wi-Fi, but speeds are still getting faster. The second group of 802.11 standards are the modern ones whose data rates exceed 1 Gbps.

Remember that Exam Objective 2.2 tells you to be familiar with 802.11 standards. Be sure to understand the differences between the newer ones.

802.11ac

In January 2014, *802.11ac* was approved, and you will often see it marketed as *Wi-Fi 5*. In many ways, it's a more powerful version of 802.11n in that it carries over many of the same features while adding in only a few new ones. It's the first commercial wireless standard that claims to offer the speed of Gigabit Ethernet.

802.11n introduced channel bonding and MIMO, and 802.11ac takes those concepts further. Instead of bonding two channels, 802.11ac can bond up to eight for a 160 MHz bandwidth. This results in a 333-percent speed increase. And 802.11ac greatly enhances MIMO. First, it doubles the MIMO capabilities of 802.11n to eight streams, resulting in another 100 percent speed increase. Second, it introduces multiuser MIMO (MU-MIMO) for up to four clients. MU-MIMO allows multiple users to use multiple antennae for communication simultaneously, whereas MIMO only allowed for one such connection at a time on a device.

The theoretical physical speed of 802.11ac is 6.9 Gbps, but most 802.11ac devices peak at about 1.3 Gbps. Common maximum throughput is just under Gigabit Ethernet speeds, at around 800 Mbps. You might see devices in the marketplace that claim to offer speeds over

2 Gbps, but the reality is that you're unlikely to get those speeds in anything less than pristine, laboratory-like conditions with all top-of-the-line hardware. In other words, it's fast, but don't count on it being *that* fast.

Back to DFS

Remember that Wi-Fi installations using the 5 GHz range need to steer clear of radar signals to avoid conflicts. Radar for airplanes and weather stations has priority over your Wi-Fi network. (Sorry!)

802.11ac obtains its fast performance mostly through channel bonding. If you will recall, 802.11n can bond 20 MHz channels into 40 MHz channels, whereas 802.11ac can take the same channels and bond them further into either 80 MHz or 160 MHz channels.

Ignoring DFS for a moment, a maximum of six nonoverlapping 80 MHz channels and two nonoverlapping 160 MHz channels are available in the 5 GHz spectrum. You can't ignore DFS, though, and it takes the maximum number of nonoverlapping 80 MHz channels down to two and eliminates any possible 160 MHz channels.

Why is this important to know? Well, mostly because it explains why it's hard to get gigabit speeds out of 802.11ac. And if a network is using only the two nonoverlapping 80 MHz channels, it will be limited to a small physical area. This makes it challenging (if not impossible) to deploy in anything other than a relatively small office. The other option is to use 40 MHz channels just like 802.11n, but then the performance will be barely better than 802.11n. If you look back at Figure 9.3, it illustrates the available channels in the 5 GHz frequency. The channels in the UNII-1 and UNII-3 are the ones that are completely available for Wi-Fi network use. UNII-2 and UNII-2 Extended channels are the DFS ones. (UNII stands for Unlicensed National Information Infrastructure and is sometimes abbreviated as U-NII.)

The most important new feature of 802.11ac is *beamforming*, which can allow for range increases by sending the wireless signal in the specific direction of the client, as opposed to broadcasting it omnidirectionally. Beamforming helps overcome the fact that the range for a 5 GHz signal is inherently shorter than one for a 2.4 GHz signal. Not all 802.11ac routers support beamforming, however, so you might have some range limitations, depending on your hardware. And even if the router does support the technology, the maximum distance still won't be any more than what you will get out of 802.11n.

802.11ax

The next version of Wi-Fi was released in 2019 and is known as *Wi-Fi 6*, with the technical designation of *802.11ax*. It gives network users what they crave—more speed. It also allows for more simultaneous users of any given access point, which is a big bonus as well.

Playing the Name Game

Wi-Fi 6 is purely a marketing name—it has nothing to do with speed or any other technical improvements. Several years ago, the Wi-Fi Alliance decided that the common public would be confused with the existing Wi-Fi naming conventions, so they needed to make it easier to understand. At the time, 802.11n was given the moniker Wi-Fi 4, and 802.11ac was designated Wi-Fi 5. So of course, the next generation of Wi-Fi should be Wi-Fi 6, and that's what it is.

Speed is, of course, the major reason that newer versions of technology get produced and accepted. Wi-Fi 6 has a few other advantages over its predecessor, though. Here's a list of enhancements versus Wi-Fi 5:

Faster—Much Faster Let's get the most user-noticeable one out of the way first. The theoretical maximum speed of Wi-Fi 6 goes from 6.9 Gbps to 9.6 Gbps. Earlier we said that even with Wi-Fi 5's theoretical maximum, in practice it was hard to get anything above 1 Gbps. With Wi-Fi 6 it's very possible to breach that threshold, albeit at pretty short distances (say, 50 feet or less).

Better Connection Management Wi-Fi 6 introduces a new modulation technique called *Orthogonal Frequency Division Multiple Access (OFDMA)*, which is an enhancement over the previously used OFDM. While OFDM was fast, it had a limitation that it could only transmit to one recipient at a time. OFDMA can handle communications with several clients at once.

Wi-Fi 6 also improved the modulation scheme from 256-QAM to 1024-QAM, going from 8 bits per symbol to 10 bits per symbol.

In addition, the Wi-Fi 5 implementation of MU-MIMO was only for downlink connections—the router could send signals to multiple receivers. With Wi-Fi 6, MU-MIMO works for uplink connections, too, meaning the router can also simultaneously receive data from multiple clients at once. This lowers latency (time spent waiting) and allows for more simultaneous devices on one network.

Less Co-Channel Interference As we noted with Wi-Fi 5, the channel bonding needed to achieve higher speeds was problematic in that it severely limited the number of non-overlapping channels available. A network with multiple wireless access points should have their coverage ranges overlap so users never hit Wi-Fi dead spots. The thing is, those ranges need to be on different, nonoverlapping channels to avoid interference. If there are only one or two channel options, then there's a problem.

In Wi-Fi 6, a feature called *Basic Service Set (BSS) coloring* adds a field to the wireless frame that distinguishes it from others, reducing the problems of co-channel interference. Specifically, an 802.11ax access point has the ability to change its color (and the

color of associated clients) if it detects a channel conflict with another access point on the same channel. If you refer to Figure 9.2, networks using Wi-Fi 6 or newer should never have to worry about interference areas because of this. It's a very cool feature sure to be underappreciated by network users, but not network administrators.

Simultaneous Use of Frequencies Wi-Fi 6 can operate at the 2.4 GHz and 5 GHz frequencies at the same time, which also increases performance. Each client device can connect on only one frequency at a time, however. Wi-Fi 6E also introduces the use of the 6 GHz frequency. It will deliver more throughput but only at very short range.

Improved Battery Life and Power Consumption for Connected Devices Finally, Wi-Fi 6 has improvements in the technology that allow for client devices to essentially "sleep" while not directly communicating with the access point. This improves battery life and reduces power consumption for client Wi-Fi devices.

In 2020, an enhancement to Wi-Fi 6 was released in the United States, called *Wi-Fi 6E*. There big difference between the two is that Wi-Fi 6E supports the 6 GHz band as well as 2.4 GHz and 5 GHz. Otherwise, they are essentially identical, based on the 802.11ax standards.

802.11be

The latest and greatest Wi-Fi standard is currently *Wi-Fi 7*, or *802.11be*. Although it wasn't ratified until early 2025, Wi-Fi 7 devices started appearing on the market in 2024.

With other technology standards, such as PCIe, new devices don't start hitting the market until well after the standards are ratified. That isn't the case with Wi-Fi. The reason is that the IEEE will publish draft guidelines about a year before they are finally ratified, and they generally don't change much or at all during that time. So, once the final draft is released (January 2024, in this case), the Wi-Fi alliance announces its certification guidelines and device manufacturers get a jump on getting the new devices into the market.

Wi-Fi 7 provides a huge speed upgrade—46 Gbps PHY—thanks to the use of 320 MHz channels, 4096-QAM modulation, and 16 spatial streams. Once again, don't expect real-life speeds to be close to that fast. Here are the details of several game-changing features:

Speed As mentioned, Wi-Fi 7 has a PHY rate of 46 Gbps, which is nearly five times faster than Wi-Fi 6. It achieves this with two key upgrades. First, it can use wider 320 MHz channels versus the 160 MHz channels that had been standard for a few generations. Second, 4096-QAM (also called 4K-QAM) modulation provides for 12 bits of data per symbol instead of 10 bits in 1024-QAM (1K-QAM), boosting speeds by 20 percent.

Multi-Link Operation (MLO) Recall that Wi-Fi 6 was the first that allowed wireless access points and routers to connect to devices on multiple frequencies at one time. Wi-Fi 7 takes it a step further and allows a single device to connect to the router across multiple frequencies simultaneously. This is great because it means more potential bandwidth and faster speeds, reduced latency, and increased connection reliability. For example, if a client device moves out of 6 GHz signal range, it will seamlessly transition over to a lower frequency and maintain the connection.

Multi-AP Coordination The short explanation of this feature is that wireless routers and access points can now talk to each other and coordinate efforts. Again, this is an enhancement of Wi-Fi 6's ability to use BSS coloring to avoid channel interference, only taken to the next level through the use of automated frequency coordination (AFC). The Multi-AP feature, combined with MLO, also allows clients to connect to multiple access points at the same time.

Better Stream Support Wi-Fi 7 supports both downlink and uplink MU-MIMO and doubles the spatial streams from 8 to 16.

Preamble Puncturing This feature is also called punctured transmission. It allows a device to isolate parts of a channel that are experiencing interference and section them off, essentially ignoring the interference.

All of these features help Wi-Fi 7 support advanced capabilities such as Augmented Reality/Virtual Reality (AR/VR) and 8K video streaming. Table 9.3 compares and contrasts the newer Wi-Fi standards. You might notice that in the table, we didn't call out the maximum distance for these versions. Each of them will have a maximum range of about 30 meters indoors and 120 meters outdoors.

TABLE 9.3 Gigabit 802.11 standards

Version	Wi-Fi 5	Wi-Fi 6	Wi-Fi 6E	Wi-Fi 7
Year	2013	2019	2020	2024
Standard	802.11ac	802.11ax	802.11ax	802.11be
Frequencies	5 GHz	2.4, 5 GHz	2.4, 5, 6 GHz	2.4, 5, 6 GHz
PHY speed	6.9 Gbps	9.6 Gbps	9.6 Gbps	46 Gbps
Max Channel size	160 MHz	160 MHz	160 MHz	320 MHz
Spatial streams	4	8	8	16
Modulation	OFDM 256-QAM	OFDMA 1024-QAM	OFDMA 1024-QAM	OFDMA 4096-QAM
MU-MIMO	DL only	UL/DL	UL/DL	UL/DL
Security	WPA2	WPA3	WPA3	WPA3
New features	Beamforming 160 MHz channel MU-MIMO	OFDMA 1024-QAM BSS coloring WPA3	6 GHz	320 MHz channel 4096-QAM MLO Multi-AP Preamble puncturing

Wi-Fi 8 Is Coming

It shouldn't shock you to know that talk about Wi-Fi 8's development was hot even before Wi-Fi 7 was officially ratified by the IEEE. Wi-Fi 8, officially known as IEEE 802.11bn Ultra High Reliability (UHR), is expected to be ratified in 2028. Instead of another speed upgrade, Wi-Fi 8 is expected to keep Wi-Fi 7's core features and focus on reliability enhancements.

While Wi-Fi 8 may not increase the PHY rate, it's expected to increase the real data rate, hopefully doubling it. It will also work to enhance the new Multi-AP and MLO features to improve seamless connectivity and reduce the latency experienced by client devices.

802.11 Devices

If you think about a standard wired network and the devices required to make the network work, you can easily determine what types of devices are needed for 802.11 networks. Just as you do on a wired network, you need a wireless network card and some sort of central connectivity device.

Wireless network cards come in a variety of shapes and sizes, including USB, PCIe, and PCI. As for connectivity devices, the most common are wireless routers (as shown in 9.4) and a type of switch called a *wireless access point (WAP)*. WAPs look nearly identical to wireless routers and provide central connectivity like wireless routers, but they don't have nearly as many features. For example, WAPs don't have the ability to act as a DHCP server or manage network security. Again, think of a WAP as an analog to a wired switch—it's there to connect client devices to the network.

FIGURE 9.4 Wireless router

Most wireless routers and WAPs also have wired ports for RJ45 connectors. The router shown in Figure 9.4 has four wired connections—in the figure you can't see them all, but they're available. The connected cable in this example is plugged into the port labeled Internet, which in this case goes to the DSL modem providing Internet access. We'll talk much more about installing and configuring a wireless router in the "Configuring Wireless Routers and Access Points" section later in this chapter.

Bluetooth Networking

We introduced Bluetooth in Volume 1 Chapter 6, "Networking Fundamentals," in the context of a wireless personal area network (PAN). It was released in 1998 and is an industry standard, much like 802.11. However, Bluetooth is not designed to be a WLAN and therefore does not directly compete with Wi-Fi. In other words, it's not the right technology to use if you want to set up a wireless network for your office. It is, however, a great technology to use if you have wireless devices that need to communicate with each other over short distances. Examples include smartphones, headsets, mice, keyboards, and printers.

Bluetooth is included in Exam Objective 2.2. Be able to explain features and uses of the technology.

Nearly every laptop comes with built-in Wi-Fi capabilities, and most also come Bluetooth-enabled. If not, install a USB Bluetooth adapter to have it communicate with Bluetooth devices. It's a safe bet to say that all smartphones and other mobile devices today support Bluetooth.

Like several versions of Wi-Fi, Bluetooth uses the unlicensed 2.4 GHz range for communication. To avoid interference, Bluetooth can "signal hop" at different frequencies to avoid conflicts with devices using other technologies in the area. Classic Bluetooth uses the FHSS modulation technique whereas Bluetooth low energy uses DSSS. Both employ a modulation scheme called Gaussian frequency shift keying (GFSK), which is a big part of the reason Bluetooth and Wi-Fi are not compatible. Thanks to technology improvements, interference with Wi-Fi is unlikely, but it can still occur.

Bluetooth Versions

Several Bluetooth standards have been introduced. Newer versions have increased speed and compatibility with technologies such as Wi-Fi, IPv6, and Internet of Things (IoT) devices, along with reduced power requirements and increased security. The newest version is Bluetooth v5.4, which was introduced in 2023. Table 9.4 provides a high-level overview of the major versions and some key features.

TABLE 9.4 Bluetooth major versions and features

Version	Basic rate (BR)	Enhanced data rate (EDR)	High speed (HS)	Low energy (LE)	Slot availability masking (SAM)
1.x	X				
2.x	X	X			
3.x	X	X	X		
4.x	X	X	X	X	
5.x	X	X	X	X	X

Now let's talk about what some of these features mean:

Basic Rate (BR) All Bluetooth devices since version 1.0 have supported a basic data transfer rate of 1.0 Mbps. After overhead, actual data transfer is actually 721 Kbps. That is incredibly slow today, but it was sufficient for early keyboards, mice, and speakers.

Enhanced Data Rate (EDR) Version 2.0 supported EDR, which delivered bit throughput of 3.0 Mbps. After overhead, it goes down to 2.1 Mbps. EDR consumes more power than BR consumes.

High Speed (HS) With version 3.0 came the HS enhancement, which allowed Bluetooth to transmit over available 802.11 signals. This boosted the data transfer rate to 24 Mbps. Keep in mind, of course, that this requires Wi-Fi to be nearby, and using HS requires considerably more power than EDR.

Low Energy (LE) Version 4.0 introduced a low energy mode, sometimes called Bluetooth Low Energy (BLE), which sacrifices data transfer rate (270 Kbps) but maintains transmission distance while using less energy. Devices using LE will consume somewhere between 1 percent and 50 percent of power compared to a device operating in Classic mode. The primary intent was to make this compatible with IoT devices. Version 4.2 was the first to include specific features for IoT.

Bluetooth devices operate in either Classic mode (BR/EDR) or low energy (LE) mode. Classic modes have higher data throughput but require more power. LE sacrifices throughput to save power. Computer peripherals are more likely to use Classic mode, whereas IoT devices, Bluetooth beacons, and things like activity monitors will probably use LE.

Slot Availability Masking (SAM) One potential issue with Bluetooth is interference on its communication channels—it uses the same unlicensed 2.4 GHz frequency that Wi-Fi does—and SAM is designed to help reduce problems. SAM can detect and prevent interference by switching bands within that frequency spectrum to help maintain throughput.

Bluetooth v5 has several new features over its predecessor, v4.2. Along with introducing SAM, and better security, it is capable of doubling the throughput and achieving four times the maximum distance, up to about 240 meters (800 feet) outdoors with line-of-sight, when in LE mode. That drops to about 40 meters (133 feet) indoors. (Remember, when distances are stated, that's the theoretical maximum under ideal conditions.) It can't do both at once, though. It can increase throughput at a shorter distance, or it can go up to longer distances at a lower data rate. It's the first version that truly challenges other IoT technologies in that space. Subsequent improvements on v5 (v5.1, v5.2, v5.3, and v5.4) have added features that locate and track devices, improve caching, upgrade LE power control and LE audio, enhance encryption, and extend battery life. All versions of Bluetooth are backward compatible with older versions. Of course, when using mixed versions, the maximum speed will be that of the older device.

One of the key features of Bluetooth networks is their temporary nature. With Wi-Fi, you need a central communication point, such as a WAP or router. Bluetooth networks are formed on an *ad hoc* basis, meaning that whenever two Bluetooth devices get close enough to each other, they can communicate directly with each other. This dynamically created network is called a *piconet*. Bluetooth-enabled devices can communicate with up to seven other devices in one piconet. One device will be the primary, and the others will be secondaries. The primary controls communication between the devices. Multiple piconets can be combined to form a *scatternet*, and it's possible for a primary of one piconet to be a secondary in another. In a scatternet, one of the Bluetooth devices serves as a bridge between the piconets.

Bluetooth devices can be configured to have an effectively permanent connection when they are within range of each other. This is convenient for smartphones and hands-free headsets or phones and automobile communications systems. To do this, the devices need to be paired. We will cover the process of pairing Bluetooth devices in Volume 1 Chapter 10, "Mobile Connectivity and Application Support."

Bluetooth Devices

As mentioned earlier, Bluetooth devices have classically been computer and communications peripherals—keyboards, mice, headsets, and printers being common. Of course, smartphones and other mobile devices support Bluetooth as well. With the newest versions, the number of IoT devices with Bluetooth capabilities has exploded as well.

One such device is a Bluetooth beacon, which is a small hardware transmitter that uses Bluetooth LE. It's broadcast only and transmits its unique identifier to nearby Bluetooth-enabled devices, such as smartphones. It can be used to send information such as marketing materials or coupons to someone with a smartphone in the vicinity of a product, or as a short-range navigation system.

There are four classes of Bluetooth devices, which differ in their maximum transmission range and power usage; the specifications are shown in Table 9.5. Most computer peripheral Bluetooth devices are Class 2 devices, which have a range of 10 meters (33 feet) and power usage of 2.5 mW. Most headsets are also Class 2, but some Class 1 headsets exist as well. Right now you might be confused, recalling from the standards discussion that Bluetooth

v5 has a maximum range of 240 meters outdoors. That is for a Class 1 device running in LE mode only; devices running in classic BR or EDR modes will have shorter ranges.

When studying Bluetooth, it's easy to get the versions and classes mixed up. Remember that they are independent of each other. Most devices you see will be Class 2, regardless of the Bluetooth version they support.

TABLE 9.5 Bluetooth device classes and specifications

Class	Distance	Power usage
1	100 meters	100 mW
2	10 meters	2.5 mW
3	1 meter	1 mW
4	0.5 meters	0.5 mW

Radio Frequency Networking Standards

The final group of networking standards we will examine is radio frequency. Technically speaking, all of the networking technologies we've discussed so far in this chapter use radio frequencies to communicate, so perhaps we're taking a bit of creative license here. The two technologies in this section are radio frequency identification and a subset of it called near-field communication.

Radio frequency identification (RFID) and near-field communication (NFC) are included in Exam Objective 2.2. Understand how they work and how they are used in networking.

Radio Frequency Identification

Radio frequency identification (RFID) is a communications standard that uses radio waves to facilitate communication. There are three types of RFID, based on the frequency used. This also affects the maximum distance that the waves can travel. Table 9.6 shows the three versions of RFID.

TABLE 9.6 RFID frequencies and characteristics

Name	Frequency	Distance
Low frequency (LF)	125–134 kHz	10 centimeters
High frequency (HF)	13.56 MHz	30 centimeters
Ultra-high frequency (UHF)	856–960 MHz	100 meters

The primary purpose of RFID is to identify items. Those items can be inventory in a store or warehouse, people, or even fast-moving things, such as race cars. An RFID system is made of three components: tag, reader, and antenna. Let's discuss what each component does:

Tag An RFID tag is fastened to the item that needs to be tracked. This can be temporary, such as an access badge an employee carries around, or it can be permanently affixed to an item. The RFID tag contains identifying information, such as an employee ID, product number, inventory number, or the like.

There are passive RFID tags and active RFID tags. Passive tags do not have a power source and draw their power from radio waves emitted by the RFID reader. This works only across short distances, typically about 25 meters or less. An active tag has its own power source (often a small battery) and may have its own antenna as well. Because it has power to generate a signal, the range for active tags is about 100 meters.

Reader The reader's job is to read information from tags. Readers have a power source and antenna. Once they get within 100 meters of an active tag or 25 meters of a passive tag, they can detect its presence. Readers can be mobile, such as a handheld inventory tracker, or static. An example of a static reader is the security gates that line the entrances of many retail stores. The products in the store have an RFID tag, and if an active tag passes through the scanners, an alarm sounds to notify the store personnel that someone may be attempting to steal an item.

Antenna An antenna boosts the distance an RFID signal can travel. Readers may have multiple antennae, depending on the model.

RFID is simplistic in networking terms. Its function is to identify items within a relatively short range. Two-way communication is pretty limited.

Near-Field Communication

A subset of RFID is a very short distance technology known as *near-field communication (NFC)*. NFC is designed to facilitate information sharing and, in particular, contactless payment. It transmits at 13.56 MHz, which is the same frequency as HF RFID.

The field of mobile contactless payment has made NFC explode in popularity over the last several years. Since Apple's introduction of the iPhone 6 back in 2014, nearly every smartphone manufacturer today equips its phones with NFC. Many tablets have NFC as well. Apple got into the NFC payment arena in 2014 with the launch of Apple Pay, which can be used from iPhones, iPads, and the Apple Watch. In 2015, Google introduced Android Pay, which is an update to the older Google Wallet app. No longer do users need to carry around credit cards—their phone and their fingerprint are all they need to complete a purchase.

NFC uses radio frequency (RF) signals, and NFC devices can operate in three different modes:

NFC Card Emulation Mode This lets the device act as a smartcard. This is useful for making payments at the site of a merchant that uses NFC.

NFC Reader/Writer Mode This allows the device to read information stored in an NFC tag on a label or poster.

NFC Peer-to-Peer Mode This allows for ad hoc data transfer between two NFC-enabled devices.

Data rates are rather slow compared to other wireless methods, as NFC operates at 106 Kbps, 212 Kbps, and 424 Kbps. NFC always involves an initiator and a target. Let's say that you wanted to read an NFC tag in a poster. You would move your phone close to the tag, and the phone would generate a small RF field that would power the target. Data could then be read from the tag. Tags currently hold up to about 8 KB of data, which is more than enough to store a URL, phone number, or other date and time or contact information.

In peer-to-peer mode, NFC data is transmitted in the NFC Data Exchange Format (NDEF), using the Simple NDEF Exchange Protocol (SNEP). SNEP uses the Layer 2 Logical Link Control Protocol (LLCP), which is connection-based, to provide reliable data delivery.

To use NFC, a user simply moves their device within range (about 10 centimeters, or 4″) of another NFC-enabled device. Then, using an app, the device will be able to perform the desired transaction, such as making a payment, reading information, or transferring data from one device to another.

NFC uses two different coding mechanisms to send data. At the 106 Kbps speed, it uses a modified Miller coding (delay encoding) scheme, whereas at faster speeds it uses Manchester coding (phase encoding). Neither method is encrypted, so it is possible to hack NFC communications using on-path or relay attacks. (We'll go into detail about specific types of attacks in Volume 2 Chapter 5, "Security Concepts.") Because of the limited distance of the RF signals, though, hacking is pretty hard to do. The potential attacker would need to be within a meter or so to attempt it.

Because of the popularity of mobile payments, it's likely that NFC will be around for quite a few years.

Installing and Configuring SOHO Networks

You already know that for computers to talk to each other, they need to be connected in some way. This can be with physical wires or through the air with one of several wireless technologies. The type of connection you choose depends on the purpose of the connection and the needs of the user or users.

Nearly every small office has a network, and it seems like most homes these days have one or more computers that need access to the Internet. As a technician, you may be asked to set up or troubleshoot any number of these types of networks, often collectively referred to as *small office, home office (SOHO) networks*. This part of the chapter will give you the background you need to feel comfortable that you can get the job done. Most of the principles we talk about apply to larger networks as well, so they're helpful if you're in a corporate environment, too.

Before we get into installation and configuration, though, it's critical to introduce a topic that permeates this whole discussion: *planning*. Before installing a single device or making changes to an existing network, *always* plan ahead. When planning ahead, consider the user's or organization's needs for today and the future. There is no sense in going overboard and recommending a top-of-the-line expensive solution if it's not needed, but if the network is likely to expand, a little up-front planning can save a lot of money and hassle in the long run.

In the following sections, we will look at how to plan and set up a SOHO network. Be advised that much of what we'll discuss from here on out isn't specifically listed as an A+ exam objective. However, we encourage you to read and understand it for two reasons. One, it has tangible real-life implications if you plan on working with computers. Two, even though the subject matter might not be a specific test objective, it will be related to subjects that are. For example, we'll touch on things like wireless channels when talking about setting up routers, how firewalls work, and differences between cable types, all of which are included in the official exam objectives. In other words, think of this next section as a combination of new material, connecting the dots on concepts you've already learned, and real-world application. With that, let's start by planning a network.

Keys to Planning a Network

Before you run your first cable or place your first wireless router, know exactly where everything is supposed to go on the network. The only way you'll be able to do this is to plan ahead. If you have planned the installation before you begin, the actual physical work of installing the network will be much easier.

Every network is going to be somewhat different. If you are installing a home-based network, the planning is usually pretty simple: figure out where the Internet connection comes

in, set up a wireless router, and configure wireless devices, such as laptops, smartphones, and home automation devices, to get on the network. If the network is going to be more complex, however, you should keep the following things in mind as you go through the planning process:

Understand relevant regulations. If you are installing a home office network, this is probably not needed. But if you are installing a network in an office building, there are likely local building codes that need to be followed.

Make a map. Understand the layout of the space in which you're installing the network. Get a map of the office or draw one yourself. Figure 9.5 shows what an office network layout could look like—there are many software packages available to make it look more professional. Add distances or a scale, if possible, so that you can determine how far you'll need to run cables or how many switches or wireless access points you'll need. Label power locations and their capacity. Mark any potential obstacles or hazards that you may run into when you try to run cable, such as fluorescent lights, water pipes, or cinder block walls.

Locate the server(s). When installing a small network, you may not have to worry about this. But if you have a network with one or more dedicated servers, decide where they will be located. They need to be in a secured location where only authorized people have access to them. This can be anything from a small closet to an elaborate server room with raised antistatic floors. Just make sure that it's temperature-controlled because server closets tend to get very hot, and we know that heat and computers don't mix well.

Identify where client computers will be. If you are setting up an office in a cubicle farm, just assume one computer (or more, depending on the cubicle type) per cube. This will help you determine where you need shared network resources as well as cable placement. (Even if all clients are running wireless, you will still need to run cables to the wireless access points.)

Locate network resources. Network users may need to share resources such as printers, so where will they be located? If there are dozens or even hundreds of users, you may need multiple printer locations or *printer banks*. Locate these and other shared resources in enough places so that users don't have to walk from one end of the office to the other just to pick up printouts.

Determine how users will connect. If network users will all connect wirelessly, you can start figuring out how many wireless routers or access points you'll need. The best way to do this is to perform a wireless site survey. The rule of thumb for Wi-Fi 5 and older is no more than 30 users per access point. Wi-Fi 6 and newer can handle more, but you will still cause performance issues if you cram too many people into one access point.

If you are going to have wired connections, start determining how long the cable runs will be. Remember that UTP has a maximum segment distance of 100 meters. If you have to go up from a patch panel, into a ceiling, and down through a wall or conduit, take that into account, too.

Designate additional connectivity areas, if needed. If you are running cables and some systems are outside of your maximum cable length, you will need to install a repeater of some sort. The best choice is probably a switch, which repeats signals. If you have several hundred computers, though, and you want to separate out networks, then a router is the best choice. These connectivity locations can be just a small closet. Other times, if no space is available, some administrators will put the switch in the drop ceiling. Although there is nothing wrong with this (as long as it's secured), it can be challenging to find power up there, and it does make it more difficult to add to that switch. Finally, if there's no way to run power into the area where you need the switch, you could buy one that uses *Power over Ethernet (PoE)*, which is covered in Volume 1 Chapter 6. Generally the number of ports these devices support is limited, but it beats having no connectivity at all.

FIGURE 9.5 Office network map

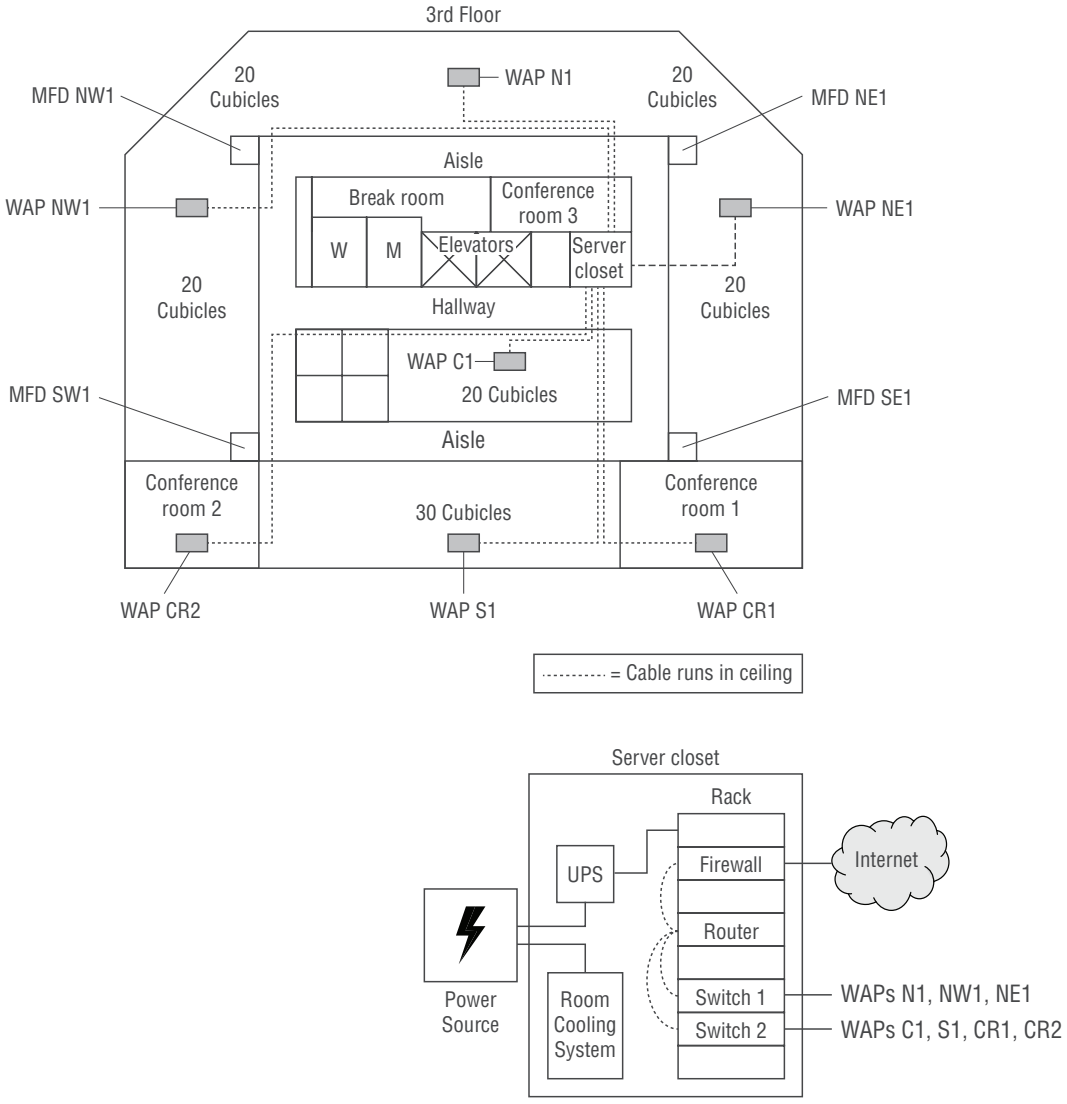

Choosing an Internet Connection

While it may be apparent to you, your clients might not realize that in order to get on the Internet, computers need an Internet connection. Your Internet connection will give you online service through an *Internet service provider (ISP)*. The type of service you choose will be determined by a few factors, such as:

- **Local provider capabilities:** Some areas have fiber-optic and copper broadband options available, but fiber isn't everywhere yet.
- **Budget considerations:** Fast Internet connections are great, but more costly than slower ones.
- **Location challenges:** For a remote mountain cabin or an island dream home, wired options might not exist.

We'll outline some of the features of each type of service and discuss why you might or might not recommend a specific connection type based on the situation.

DSL

For a couple of decades, copper-based broadband was the king of high-speed Internet access. While it's fading in popularity to even faster fiber-optic choices, it's still used today. One of the two most popular copper wire-based broadband choices for home use is *digital subscriber line (DSL)*. It uses existing phone lines and provides fairly reliable high-speed access. To utilize DSL, you need a DSL modem (the back of which is shown in Figure 9.6) and a network card in your computer. The ISP usually provides the DSL modem, but you can also purchase them in a variety of electronics stores. You use an Ethernet cable with an RJ45 connector to plug your network card into the DSL modem and the phone cord to plug the DSL modem into the phone outlet. If you need to plug a landline into the same phone jack as your DSL modem, you will need a DSL splitter (such as the one shown in Figure 9.7) and plug the splitter into the wall.

Instead of plugging your computer directly into the DSL modem, you can plug your computer into a router (such as a wireless router) and then plug the router into the DSL modem. This is a great way to share a broadband connection with multiple devices.

There are several different forms of DSL, including *high bit-rate DSL (HDSL)*, *symmetric DSL (SDSL)*, *very high bit-rate DSL (VDSL)*, and *asymmetric DSL (ADSL)*. Each of those standards has different maximum speeds, but you won't need to know them for the A+ exam, and probably not in real life, either. Some key things to know are:

- Most DSL subscriptions are asymmetric, meaning that download speed will be much faster than upload speed.
- Generally speaking, DSL tops out at about 70 Mbps download and 20 Mbps upload.

FIGURE 9.6 A DSL modem

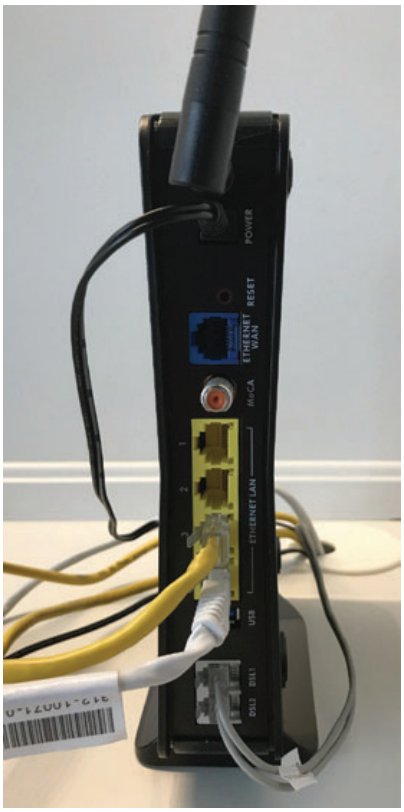

FIGURE 9.7 A DSL splitter

- The maximum speeds decrease as the installation gets farther away from the phone company's equipment.
- When contacting a phone company to get DSL, you will be quoted a price for a speed, and the technology behind it (VDSL versus ADSL, for example) is sort of immaterial. At most, you *might* get options for symmetric versus asymmetric plans.
- It's uncommon to get the maximum speed quoted by providers for a given plan.
- Because DSL uses copper telephone lines, the DSL modem connects to the wall using an RJ11 connector.

DSL is part of Exam Objective 2.7, "Compare and contrast internet connection types, network types, and their characteristics." Understand how DSL works (over phone lines), approximate maximum speeds, and how to connect DSL for SOHO use.

ADSL was the most popular form of DSL for many years, and intuitively it makes a lot of sense based on how users use the Internet. Most Internet traffic is downloaded, not uploaded. Imagine a 10-lane highway. If you knew that 8 out of 10 cars that drove the highway went south, wouldn't you make eight lanes southbound and only two lanes northbound? That is essentially what ADSL does.

DSL and your voice communications can work at the same time over the phone line because they use different frequencies on the same wire. Regular phone communications use frequencies from 0 to 4 kHz, whereas DSL uses frequencies in the 25.875 kHz to 138 kHz range for upstream traffic and in the 138 kHz to 1,104 kHz range for downstream traffic. Figure 9.8 illustrates this.

FIGURE 9.8 Voice telephone and ADSL frequencies used

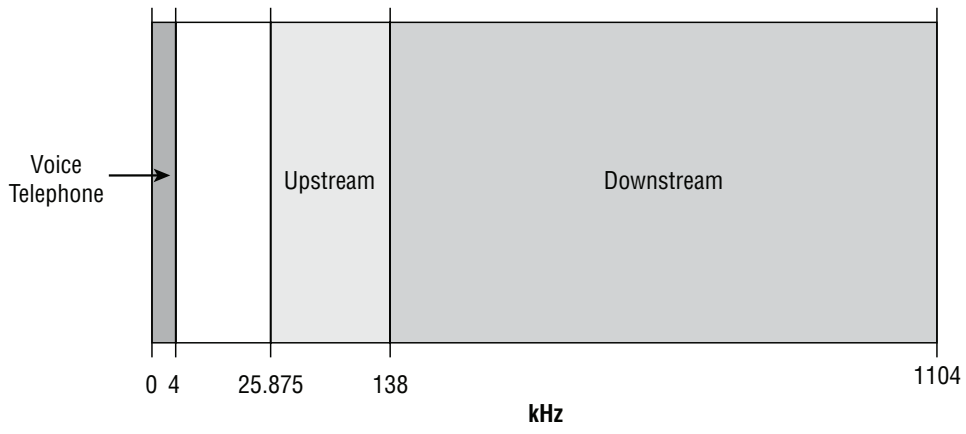

One major advantage that DSL providers tout is that with DSL you do not share bandwidth with other customers, whereas that may not be true with cable modems.

You will probably see Internet providers offering broadband speeds of 1 Gbps or more, and they might call it DSL. Technically, though, speeds that fast are fiber-optic connections, which we cover later in the section "Fiber-Optic Internet." Over the several years, the lines between DSL and fiber-optic have blurred considerably. It used to be that you ordered one or the other. Today, it's more common to order a broadband speed, and the provider figures out if they will give it to you via traditional copper wires (DSL) or fiber-optic.

DSL remains a popular choice for both small businesses and residential offices. If it's available, it's easy to get the phone company to bundle service with a landline and bill both at the same time. Often there will be a package discount for having multiple services.

With many people using their smartphones as their home phones and landlines fading into history, you may wonder if this causes a problem for those who want DSL. Not really. Most phone providers will provide you with DSL without a landline (called *naked DSL*). Of course, you are going to have to pay a surcharge for the use of the phone lines if you don't already use one.

Mr. Watson, Come Here, I Need You

Those famous words were the first spoken over the telephone by Alexander Graham Bell in 1876. And 150 years later, some people still use a modem and the *plain old telephone service (POTS)* for their Internet connections. Not many, but some. In the 1990s, dial-up Internet service was the most common way to get online. Even as recently as 2000, an estimated 74 percent of U.S. residential Internet connections were dial-up. Today, it's less than one-tenth of one percent.

Most of the people who still use dial-up do it because it's cheaper than broadband or high-speed access isn't available where they live. And it's relatively easy to configure. The only hardware you need is a modem and a phone cable. You dial in to a server (such as an ISP's server), provide a username and a password, and you're on the Internet.

Of course, there are reasons a dial-up connection might not be appropriate. The big one is speed—it's limited to 56 Kbps. If people need to download files or have substantial data requirements, dial-up is probably too slow. Forget about video or audio streaming. In addition, with limited bandwidth, it's good only for one computer. It *is* possible to share a dial-up Internet connection among computers by using software tools, but it's also possible to push a stalled car up a muddy hill. Neither option sounds like much fun. If broadband isn't available in a certain location, satellite is almost certainly a better option than dial-up.

Cable

The other half of the popular copper-based, home-broadband duet is *cable Internet*. Using a cable modem, cable Internet provides high-speed Internet access through your cable service, much like DSL does over phone lines. You plug your computer into the cable modem using a standard Ethernet cable, just as you would plug into a DSL modem. The only difference is that the other connection goes into a cable TV jack (using an F-type connector) instead of the phone jack (using an RJ11 connector). Cable Internet provides broadband Internet access via a specification known as Data Over Cable Service Interface Specification (DOCSIS). Anyone who can get a cable TV connection should be able to get the service.

Cable Internet is part of Exam Objective 2.7. Be able to compare and contrast it to other Internet connection types.

As advertised, cable Internet connections are usually faster than DSL connections. You'll see a wide variety of claimed speeds; some cable companies offer packages with download speeds up to 50 Mbps, 100 Mbps, or 400 Mbps and various upload speeds as well. If it's that fast, why wouldn't everyone choose it? Although cable generally is faster, a big caveat to these speeds is that they are not guaranteed and they can vary. And again, with many phone companies not really differentiating between DSL and fiber-optic, it can be difficult to understand exactly what you're comparing.

One of the reasons that speeds may vary is that you are sharing available bandwidth within your distribution network. The size of the network varies, but it's usually between 100 and 2,000 customers. Some of them may have cable modems too, and access can be slower during peak usage times. Another reason is that cable companies make liberal use of bandwidth throttling. If you read the fine print on some of their packages that promise the fast speeds, one of the technical details is that they boost your download speed for the first 10 MB or 20 MB of a file transfer, and then they throttle your speed back down to your normal rate.

To see how this could affect everyone's speed on the shared bandwidth, consider a simplified example. Let's say that two users (Sally and John) are sharing a connection that has a maximum capacity of 40 Mbps. For the sake of argument, let's assume that they are the only two users and that they share the bandwidth equally. That would mean normally each person gets 20 Mbps of bandwidth. If Sally gets a boost that allows her to download 30 Mbps, for however long, that leaves John with only 10 Mbps of available bandwidth. If John is used to having 20 Mbps, that 10 Mbps is going to seem awfully slow.

Although it may seem as though we are down on cable Internet, you just need to understand exactly what you and your customers are getting. In practice, the speeds of a cable modem are pretty comparable to those of DSL. Both have pros and cons when it comes to reliability and speed of service, but a lot of that varies by service provider and isn't necessarily reflective of the technology. When it comes right down to it, a choice between DSL and cable, if both are available in your area, may depend on which company offers you the best deal (although realistically, the best performance choice will be fiber-optic or cellular 5G).

Cable modems can be connected directly to a computer but can also be connected to a router or wireless router just like a DSL modem. Therefore, you can share an Internet connection over a cable modem.

For detailed information about cable Internet availability and performance, check out `www.highspeedinternet.com`.

Fiber-Optic Internet

Fiber-optic cable is pretty impressive with the speed and bandwidth it delivers. For nearly all of fiber-optic cable's existence, it's been used mostly for high-speed telecommunications and network backbones. This is because it has always been more expensive than copper to install and operate. The cables themselves are pricier, as is the hardware at the end of the cables. Technology follows this inevitable path of getting cheaper the longer it exists, and fiber is following this trend. Many phone and media companies are now offering fiber-optic Internet connections for home subscribers.

Fiber Internet is part of Exam Objective 2.7. Be able to compare and contrast it to other Internet connection types.

An example of one such option is Fios by Verizon. It offers *fiber-to-the-home (FTTH)* service, which means that the cables are 100 percent fiber from their data centers to your home. As of this writing, Fios offered basic packages at 200 Mbps download and 200 Mbps upload, and the fastest speeds are 2 Gbps each direction. Other providers will offer similar packages.

Some cable companies promise a high-speed, fiber-optic connection for your TV cable as well as cable Internet service. In many cases, the connection is *fiber-to-the-node (FTTN)*, and the fiber runs only from their network to the junction box at the entrance to your neighborhood or possibly to your curb (called *fiber-to-the-curb [FTTC]*). From there, the cable is coaxial copper. This means that speeds will be significantly slower than advertised. If you're paying for a fiber connection, be sure you're actually *getting* a fiber connection.

Connecting to fiber-based Internet requires an optical network terminal (ONT), which we talked about in Volume 1 Chapter 6 and showed in Figure 6.35. From the ONT, you will have a copper network cable running to a router of some sort (say, a wireless router), and then the computers will connect to the router to get to the Internet.

Are there any downsides to a fiber Internet connection? Really only two come to mind. The first is availability. It's still a little spotty on where you can get it. The second is price. That great gigabit connection can easily cost you $200 per month after any special introductory pricing wears off.

Putting It All Together

We've covered the three wired Internet connection types in the exam objectives, and setting each of them up is relatively similar. The first step, of course, is to contact the local providers to see what's available in your area and then sign up. Oftentimes they will try to sell or lease you a wireless router (or they may call it a gateway). Our opinion is to buy your own because it's more cost effective in the long run and it gives you complete control over the device. Get only whatever Internet connection box (DSL modem, cable modem, or ONT) is required.

Likely they will have to send a technician to your location to finish the connection, possibly needing to update some cables in the process.

Once that's all completed, though, it's time to hook it up internally. We've verbally described it a few times, but also want to show a picture of what this looks like. Figure 9.9 does that.

FIGURE 9.9 Connecting DSL Internet

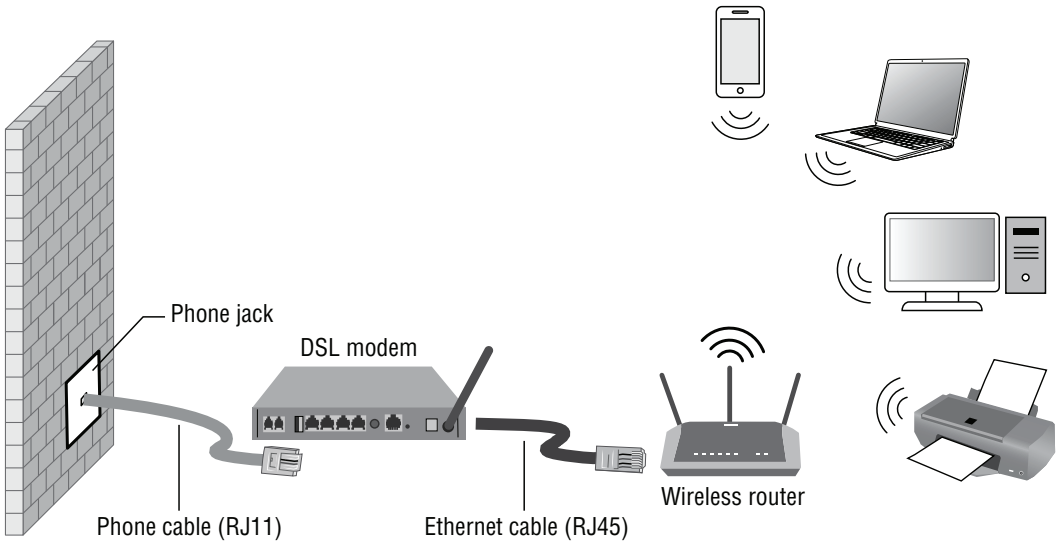

For cable Internet, the only differences are that it's a cable modem instead of a DSL modem, the wall plate will be a cable connector, and the cable used will most likely have an F-type connector.

For fiber-optic Internet, the differences are that the connectivity device will be an ONT, and the cable coming from the wall will be fiber-optic with an appropriate connector.

Satellite

Moving on from wired Internet connections, let's talk about wireless ones. One type of broadband Internet connection that does not get much fanfare is *satellite Internet*. Instead of a cabled connection, it uses a satellite dish to receive data from an orbiting satellite and relay station that is connected to the Internet. Satellite connections are typically slower than wired broadband connections, with downloads often maxing out at around 250 Mbps and uploads around 3 Mbps. To compare plans and prices available in your area, visit `https://satelliteinternet.com`.

Exam Objective 2.7 includes satellite Internet. Be able to compare and contrast it with other Internet connection types.

The need for a satellite dish and the reliance on its technology is one of the major drawbacks to satellite Internet. People who own satellite dishes will tell you that there are occasional problems due to weather and satellite alignment. You must keep the satellite dish aimed precisely at the satellite or your signal strength (and thus your connection reliability and speed) will suffer. Plus, cloudy or stormy days can cause interference with the signal, especially if there are high winds that could blow the satellite dish out of alignment. Receivers are typically small satellite dishes (like the ones used for DirecTV or Dish Network) but can also be portable satellite modems (modems the size of a briefcase) or portable satellite phones.

Another drawback to satellite technology is the delay (also called *propagation delay*), or *latency*. The delay occurs because of the length of time required to transmit the data and receive a response via the satellite. This delay (between 250 and 350 milliseconds) comes from the time it takes the data to travel the approximately 35,000 kilometers into space and return. To compare it with other types of broadband signals, cable and DSL have a delay between customer and ISP of 10 to 30 milliseconds. With standard web and email traffic, this delay, while slightly annoying, is acceptable. However, with technologies like VoIP and live Internet gaming, the delay is intolerable.

Online gamers are especially sensitive to propagation delay. They often refer to it as *ping time*. The higher the ping time (in milliseconds), the worse the response time in the game. It sometimes means the difference between winning and losing an online game.

Of course, satellite also has advantages; otherwise, no one would use it. First, satellite connections are incredibly useful when you are in an area where it's difficult or impossible to run a cable, or if your Internet access needs are mobile and cellular data rates just don't cut it, at least not until you get cellular 5G.

The second advantage is due to the nature of the connection. This type of connection is called *point-to-multipoint* because one satellite can provide a signal to a number of receivers simultaneously. It's used in a variety of applications, from telecommunications and handheld GPSs to television and radio broadcasts and a host of others.

Here are a few considerations to keep in mind regarding satellite:

It's expensive compared to other broadband access, and you might get limited data. Packages usually start around $50 per month for about 50 Mbps downloads. That kind of download speed will cost you only about $20–$30 for DSL or cable. And many providers set thresholds on the amount of data you can download per month. For an example, a popular 50 Mbps package for $50 per month limits you to 100 GB of data per month (as of this writing). Unlimited plans may be upward of $70 per month. Going over that amount can result in extra charges and/or speed throttling. Your speed will be decreased for a certain period, and you will pay more for that data as well.

Installation can be tricky. When installing a satellite system, you need to ensure that the satellite dish on the ground is pointed at precisely the right spot in the sky. This can be tricky to do if you're not trained, but some have a utility that helps you see how close you are to being right on (you're getting warmer . . . warmer).

Line-of-sight is required. Satellite communications require *line-of-sight*. A tree between you and your orbiting partner will cause problems. Rain and other atmospheric conditions can cause problems as well.

Latency can be a problem. Because of the long distance the message must travel, satellites can be subject to long latency times. While it happens with wired connections, it disproportionately affects satellite transmissions. Have you ever watched a national news channel when a reporter is reporting from some location halfway across the world? The anchor behind the desk will ask a question, and the reporter will nod, and nod, and finally about five excruciating seconds after the anchor is done, the reporter will start to answer. That's latency.

Real World Scenario

All in the Name of Entertainment

Several years ago as a teenager, one of the authors worked for a local television station during the summers. Each summer, the television station would broadcast a Senior PGA golf tournament that was held on a nearby mountain course.

Before the tournament, the crew would spend three days setting up the control truck, cameras, and link back to the station. (It was a network with TV cameras instead of workstations.) Because of the remote location, the crew had to set up a satellite uplink to get the signals back to civilization. From the control truck, a transmitter was pointed at a relay station on the side of the mountain, which in turn was pointed at a satellite orbiting the earth. It took a team of four engineers to set it up. Two engineers would stay at the truck, and two others would board ATVs and journey up the remote mountainside. Once in

position, they would set up the relay station, which looked a lot like a keg of beer with a few antennas. The engineers at the truck would adjust their directional microwave transmitter until the relay station received a strong signal. Then the engineers on the mountainside would perform the arduous task of pointing their transmitter at the satellite.

It was a long and tedious process, and that's really the point of the story. Satellite was the *only* option available to complete the network, but satellite networks can be a challenge to set up and configure.

Cellular

It seems that everyone—from kindergarteners to 80-year-old grandparents—has a smartphone today, and of course almost all of them have persistent Internet access. The industry has revolutionized the way we communicate and, some say, contributed to furthering an attention deficit disorder—like, instant gratification–hungry society. In fact, the line between cell phones and computers has blurred significantly with all the smartphones on the market. It used to be that the Internet was reserved for "real" computers, but now anyone can be online at almost any time.

You're probably at least somewhat familiar with cellular Internet—it's what smartphones, many tablets, and older cellular (or cell) phones use. You might not have thought of it in terms of networking, but that's really what it is. There's a central access point, like a hub, which is the cellular network tower. Devices use radio signals to communicate with the tower. The tower, in turn, is connected via wires to a telecommunications backbone, which essentially talks to all the other commercial telecom networks in the world. It's a huge network.

Cellular Internet is included in Exam Objective 2.7. Understand how it works versus other Internet connection types.

For years, this network was pretty slow, especially by today's standards. The most advanced cell standard when the Internet started becoming a thing was 3G. (There were also 1G and 2G standards before that.) Initially, it had the bandwidth to carry only voice conversations, which consume little bandwidth. Then texting was supported, and after several enhancements it could in theory support downloads of 7 Mbps, although actual data rates varied by carrier, equipment, the number of users connected to the tower, and the distance from the tower. The more current standards are 4G and 5G.

Fourth Generation (4G)

In 2008, the next generation beyond 3G, appropriately named *fourth generation (4G)*, made its appearance. To be specific, 4G refers to a generation of standards for mobile devices (such as phones and tablets) and telecommunication services that fulfill the International Mobile Telecommunications Advanced (IMT-Advanced) specifications as adopted by the

International Telecommunication Union (ITU). In more practical terms, it's simply the next-in-line standard for wireless telephone, Internet, video, and mobile TV that replaced 3G. To meet IMT-Advanced standards, the service must provide peak data rates of at least 100 Mbps for high-mobility communication (such as trains or cars) and 1 Gbps for low-mobility communication. One major difference between 4G and 3G is that 4G is designed to use IP instead of traditional telephone circuits. It's designed to provide mobile broadband access.

The first 4G devices that came on the market did not offer anything close to the speeds specified by the ITU. Mobile manufacturers branded them 4G anyway, and there wasn't much the ITU could do to stop it. The result was that the world became inundated with 4G advertising.

In the early days of 4G, there were two competing standards—WiMAX and *Long-Term Evolution (LTE)*. WiMAX was the marketing name given to the IEEE 802.16 standard for wireless MAN technology. While it was initially promising and had higher speeds than its competing standard, LTE was what the mobile providers latched onto. And as already mentioned, they advertised a lot. For years, whenever you turned on the TV, you couldn't help but be bombarded with commercials from cell providers pitching the fastest or widest or whatever-est 4G LTE network.

The biggest enhancement of 4G LTE over 3G is speed. Whereas with true 3G technology you were limited to about 500 Kbps downloads, some 4G LTE networks will give you download speeds of 10–20 Mbps and upload speeds of 3–10 Mbps. (The theoretical maximum for LTE is 300 Mbps download and 75 Mbps upload.) The range of 4G LTE depends on the tower and obstructions in the way. The optimal cell size is about 3.1 miles (5 km) in rural areas, and you can get reasonable performance for about 19 miles (30 km).

Fifth Generation (5G)

New ITU mobile specifications come out about every 10 years, so it stands to reason that we're now on 5G. Even though the first 5G modem was announced in 2016, it took until late 2018 for cellular providers to test-pilot 5G in several cities. Rollout expanded in earnest in 2019, and now it's fairly widespread, though it's not everywhere yet.

The *fifth generation (5G)* of cellular technology is a massive improvement over 4G—some users will be able to get sustained wireless speeds in excess of 1 Gbps. The theoretical maximum peak download capacity is 20 Gbps, but, of course, that would require pristine conditions, which don't occur in real life. The technical specifications for 5G divide it into three categories:

- Enhanced Mobile Broadband (eMBB), which will be used for cell phones and mobile communication
- Ultra Reliable Low Latency Communications (URLLC) for autonomous vehicles and industrial applications
- Massive Machine Type Communications (mMTC) for sensors, which is designed to support IoT devices

Initially the focus has been on eMBB and developing the infrastructure to support mobile devices for consumers. Two versions of eMBB will take hold: fixed wireless broadband in densely populated areas and LTE for everywhere else.

Let's start with LTE, because we've already talked about it some. 5G's version of LTE is similar to 4G LTE, just with faster speeds. It will use existing LTE frequencies in the 600 MHz to 6 GHz range and existing 4G LTE infrastructure. Browsing speeds for 5G are about seven to nine times faster than 4G (490 Mbps on average), and most users can get 100 Mbps download speeds, compared to 8 Mbps on their 4G LTE network. So, in general, expect 5G LTE to be about seven to ten times faster than a comparable 4G connection.

The really exciting feature of eMBB is fixed wireless broadband. This technology uses millimeter wave bands (called *mmWave*) in the 24 GHz to 86 GHz range. With mmWave, 5G users should expect gigabit speeds over a wireless connection. This great performance comes with a catch, though. (Doesn't it always?)

Very short radio waves such as the ones used in mmWave can carry a lot of data, but there are two inherent problems:

- The first is distance. They're limited to about half a mile.
- The second is obstruction. mmWave can be blocked by just about anything, including walls, trees, and even a user's hand.

To overcome the first challenge, transmitters need to be placed very close together. This shouldn't be too much of a problem in urban areas, because 4G transmitters are already packed close together and providers can simply attach a 5G transmitter in the same place.

The second challenge is a bit trickier. Engineers found a way to take advantage of signal bounce—the fact that signals blocked by buildings actually bounce off the buildings—to ultimately bounce the signal to the end user's device. It slows down transmission speeds a bit, but it's still a lot faster than 4G technology. Due to the limitations, however, the fastest 5G connections will be available only in densely populated urban areas.

The promise of 5G performance is intoxicating. With it, users can stream 8K movies over cellular connections with no delays, and it can also enable virtual reality and augmented reality over cell connections. 5G users in urban areas may never need to connect to a Wi-Fi network again, because their cellular connections will be just as fast. Wireless plans may truly become unlimited because of the bandwidth capacity that mmWave creates. For these reasons, some experts have called 5G a revolution, not an evolution.

Many cellular service providers now offer home-based Internet solutions based on 5G technology.

6G cellular is currently under development. The specifications aren't expected to be released until 2028, and it will likely be 2030 before the first 6G networks start to appear. Once 6G hits the market, 4G LTE networks will likely start to disappear.

Wireless Internet Service Provider

One final type of wireless Internet connection is that provided by a *Wireless Internet service provider (WISP)*. In very broad terms, a WISP is an ISP that grants access using a wireless technology. Specifically, though, the industry uses the term to refer to providers that offer

fixed point-to-point, relatively short-distance broadband Internet, which typically requires unobstructed line-of-sight.

Exam Objective 2.7 includes wireless Internet service provider (WISP). Understand how it works versus other Internet connection types.

That differs from the 5G we mentioned in the last section, because while 5G is point-to-point, the receiver (someone's smartphone) is not generally in a fixed location. The receiver for a WISP-based Internet connection will be a fixed receiver, often a small dish or an antenna.

WISPs can operate over unlicensed channels such as 900 MHz, 2.4 GHz, 5 GHZ, 24 GHz, and 60 GHz, or they might offer service in a licensed frequency in the 6 GHz to 80 GHz range. WISP connections require line-of-sight and can be subject to interference and delay. Table 9.7 summarizes the connection types we have discussed in this chapter.

TABLE 9.7 Common Internet connection types and speeds

Connection type	Approximate basic package cost	Download speed range	Description
DSL	$20–$30	Up to 70 Mbps	Inexpensive broadband Internet access method with wide availability, using telephone lines.
Cable	$20–$30	Up to 1+ Gbps	Inexpensive broadband Internet access method with wide availability, using cable television lines.
Fiber	$40–$50	Up to 2 Gbps	Incredibly fast and expensive.
Satellite	$40–$50	Up to 250 Mbps	Great for rural areas without cabled broadband methods. More expensive than DSL or cable.
Cellular	$30–$50	Up to 100 Mbps with 5G LTE or 1 Gbps with mmWave	Great range; supported by cell phone providers. Best for a very limited number of devices.
WISP	$40–$150	6 Mbps to 50 Mbps	Fast connection for rural areas without cabled broadband methods.

Choosing Internal Network Connections

Along with deciding how the computers will get to the outside world, you need to think about how the computers will communicate with each other on the internal network. The choices you make will depend on the speed needed, distance and security requirements, and

cost involved with installation and maintenance. It may also depend some on the abilities of the installer or administrative staff. There may be someone who is quite capable of making replacement Cat 6 cables but for whom making replacement fiber-optic cables is a much more daunting task. The choices for internal connections can be lumped into two groups: wired and wireless.

Most networks today are a hybrid of wired and wireless connections. Understand the fundamentals of how each works separately; then you can understand how they work together. Every wireless connection eventually connects back to a wired network point.

Wired Network Connections

With the prevalence of wireless networks, wired connections may seem dated. The truth is, wired connections form the backbone of nearly every network in existence. The importance of wired connections still remains strong. In general, wired networks are faster and more secure than their wireless counterparts.

When it comes to choosing a wired network connection type, you need to think about speed, distance, security, and cost. You learned about several types of wired connections in Volume 1 Chapter 6, such as coaxial, UTP, STP, and fiber-optic, but the only two you'll want to go with today are twisted pair and fiber. You'll run one of the two (or maybe a combination of the two), with UTP by far the most common choice, as an Ethernet star network.

Refer to Volume 1 Chapter 6 if you need a refresher of the various twisted pair or fiber-optic standards available. Here are things to think about when choosing between the two:

How fast does the network need to be? In today's environment, there's no reason to install anything that doesn't support gigabit speeds, so Cat 5e is the bare minimum standard to use. Realistically though, Cat 6 is basically the same price and supports 10 Gbps, which is important for future upgradeability. Go with at least Cat 6 unless fiber-optic is in the mix.

What is the maximum distance of any cable run? In most office environments, you can configure a network in such a way that 100 meters will get you from any connectivity device to the end user. Most times today, clients will be on a wireless connection in the office so the only long cable runs will be to the furthest wireless access point. If there is a run of longer than 100 meters, a repeater is needed unless you're installing fiber.

Are there any major security concerns? Copper cable is pretty secure, but it does emit a signal that can be intercepted, meaning people can tap into your transmissions (hence, the term *wiretap*). Shielded twisted pair is better than unshielded in this regard, but all copper does emit a signal. Fiber-optic cables are immune to wiretapping. Normally this isn't a big deal, because copper cables are infinitely more secure than wireless, which broadcasts your data through the air. But if security is of the utmost concern, then fiber is a better choice than copper.

Is there a lot of electrical interference in the area? Transmissions across a copper cable can be ravaged by the effects of *electromagnetic interference (EMI)* from fluorescent lights, microwave ovens, fans, generators, power lines, computers, and other devices. Fiber is immune to those effects.

How much will all of this cost? As you're thinking about which type of cable to use, also consider the hardware you'll need. If you are running fiber to the desktop, you need fiber network cards, routers, and switches. If you are running UTP, you need network cards, routers, and switches with RJ45 connectors. Fiber cables and hardware are more expensive than their copper counterparts. Table 9.8 summarizes cable choices and provides characteristics of each.

Understand that the costs shown in Table 9.8 are approximate and are for illustrative purposes only. The cost for equipment in your area may differ, and costs also decrease as technologies mature. Fiber continues to get cheaper, but it's still far more expensive than copper.

TABLE 9.8 Cable types and characteristics

Characteristics	Twisted pair	Fiber-optic
Transmission rate	Cat 5e: 1 Gbps Cat 6/6a and 7: 10 Gbps Cat 8: 25 Gbps or 40 Gbps	1 Gbps to 40 Gbps
Maximum length	100 meters (328 feet) is standard; may be shorter distances at higher speeds	About 40 kilometers or 25 miles
Flexibility	Very flexible	Fair
Ease of installation	Very easy	Difficult
Connector	RJ45	SC, ST, LC, and others
Interference (security)	Susceptible	Not susceptible
Overall cost	Inexpensive	Expensive
PCIe NIC cost	1 Gbps: $10–$40 10 Gbps: $50 and up	$70–$150; easily $300–$500 for server NICs
10-meter cable cost	Cat 6/6a: $15–$30Cat 8: $30–$50	Depends on mode and connector type, but generally $20–$50
8-port switch cost	1 Gbps: $25–$100	$300 and up

Fiber-optic cabling has some obvious advantages over copper, but, as you can see, it may be prohibitively expensive to run fiber to the desktop. What a lot of organizations do is use fiber sparingly, where it's needed the most, and then run copper to the desktop. Fiber will be used in the server room and perhaps between floors of a building as well as any place where a very long cable run is needed.

Wireless Network Connections

People love wireless networks for one major reason: convenience. Wireless connections enable a sense of freedom in users. They're not stuck to their desks; they can work from anywhere! (We're not sure if this is actually a good thing.) Wireless isn't typically as fast and it tends to be a bit more expensive than wired copper networks, but the convenience factor far outweighs the others.

When you are thinking about using wireless for network communications, the only real technology option available today is Wi-Fi. The other wireless technologies we discussed earlier in the chapter can help mobile devices communicate, but they aren't designed for full wireless LAN (WLAN) use. Your choice becomes which Wi-Fi standard you want to use.

If network users have a specific need, such as Bluetooth printing or managing IoT devices, you will want to integrate an appropriate technology for that into your design as well.

So how do you choose which one is right for your situation? You can apply the same thinking you would for a wired network in that you need to consider speed, distance, security, and cost. Generally speaking, though, with wireless it's best to start with the newest technology and work your way backward.

Security concerns about wireless networks are similar, regardless of your choice. You're broadcasting network signals through air; there will be some security concerns. And all current Wi-Fi standards have the same range, so, at the end it's speed and cost.

In today's environment, it's silly to consider 802.11n or older. Deciding that you are going to install an 802.11n network from the ground up at this point is a bit like saying you are going to use 10BaseT. You could try, but why?

That brings us to your most likely choices: Wi-Fi 7 (802.11be), Wi-Fi 6 (802.11ax), and 802.11ac (Wi-Fi 5). Wi-Fi 7 represents a pretty serious performance jump over previous versions, but the wireless routers, access points, and network cards may be more expensive for a while. It may come down to cost. Shop around to see what kind of deal you can get. Exercise 9.1 has you do just that.

EXERCISE 9.1

The Cost of Networking

1. Visit the website for an electronics store. If you're unfamiliar with any, try `www.bestbuy.com` or `www.microcenter.com`.
2. Find a Wi-Fi 7 router. How much is it?

EXERCISE 9.1 *(continued)*

3. Find an older standard. See if you can find a Wi-Fi 5 router. How much is it? (You might have to try a site like `eBay.com` to find older devices.)
4. Now price wireless network cards. How much difference is there between standards?
5. (Optional) Find a fiber-optic card and compare its price to an Ethernet card that offers similar speeds. Also look at the price of a 25-meter Cat 6 cable versus a 5-meter fiber-optic cable. How much difference is there?

Installing the Network Infrastructure

Once all your plans are complete, you've double-checked them, and they've been approved by the client or boss, you can begin physical installation. As we've said before (but can't overstate), having good plans up front saves time and money. The approval process is critical too, so the people in charge are informed and agree to the plans. Here we'll look at installation of three groups of items: network cards, cables, and connectivity devices.

Installing an Internal NIC

Before you can begin communicating on your network, you must have a NIC installed in the device. External USB NICs are super easy to install—you literally plug it in and it will install its driver and be ready to go. Installing an internal NIC is a fairly simple task if you have installed any expansion card before; a NIC is just a special type of expansion card. In Exercise 9.2, you will learn how to install an internal NIC.

EXERCISE 9.2

Installing an Internal NIC in Windows 11

Before working inside a computer case, be sure to understand and follow safety procedures, including preventing electrostatic discharge (ESD).

1. Power off and unplug the PC, remove the case and the metal or plastic blank covering the expansion slot opening, and insert the expansion card into an open slot.
2. Secure the expansion card with the screw provided.
3. Put the case back on the computer and power it up. (You can run software configuration at this step, if necessary.)

   If there are conflicts, change any parameters so that the NIC doesn't conflict with any existing hardware.

   Note that these first three steps may not be necessary if you have an onboard NIC.

4. Install a driver for the NIC for the type of operating system that you have. Windows Plug and Play (PnP) will recognize the NIC and install the driver automatically.

   Windows may also ask you to provide a copy of the necessary driver if it does not recognize the type of NIC you have installed. If Windows does not start the installation routine immediately, you can add it manually:

   a. Click Start ➢ Settings (it looks like a gear) ➢ Bluetooth & Devices.

   b. Click the plus sign that says Add Device. That will bring up the Add A Device window (Figure 9.10).

FIGURE 9.10 Add A Device window

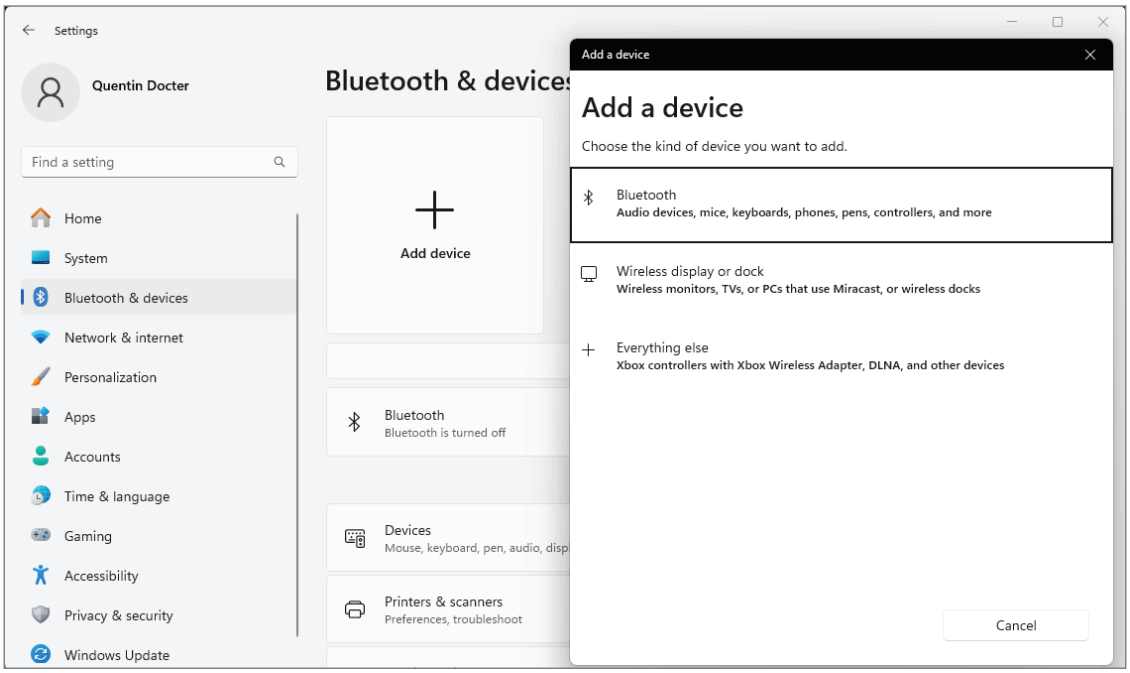

   c. Click Everything Else.

   d. When Windows finds the NIC, choose it and continue the installation.

After installing a NIC, you must hook the card to the network using the appropriate cable (if using wired connections). Attach this patch cable to the connector on the NIC and to a port in the wall (or connectivity device), thus connecting your PC to the rest of the network.

Configuring a Wired NIC

Now that your NIC is installed, it's time to configure it with the right IP address and TCP/IP configuration information. As you learned in Volume 1 Chapter 8, "Network Services," there

are two ways to do this: dynamic IP addressing from a DHCP server, or static IP addressing if it's done manually.

Bear in mind that these NIC configuration steps are general guidance. Always consult the documentation that comes with the hardware to ensure that there isn't a special step unique to that card.

To configure your NIC in Windows 11 (this is a review from Volume 1 Chapter 7, "Introduction to TCP/IP," Exercise 7.1), click Start ➢ Settings ➢ Network & Internet. Scroll to the bottom and click Advanced Network Settings. You will then see the Advanced Network Settings window, like the one shown in Figure 9.11. Find the Network adapter, and next to More Adapter Options, click Edit. That will open the Ethernet Properties window. Click Internet Protocol Version 6 (TCP/IPv6) (or Version 4 if your network is solely TCP/IPv4), then Properties. That will open the TCP/IP properties shown in Figure 9.12.

FIGURE 9.11 Advanced Network Settings

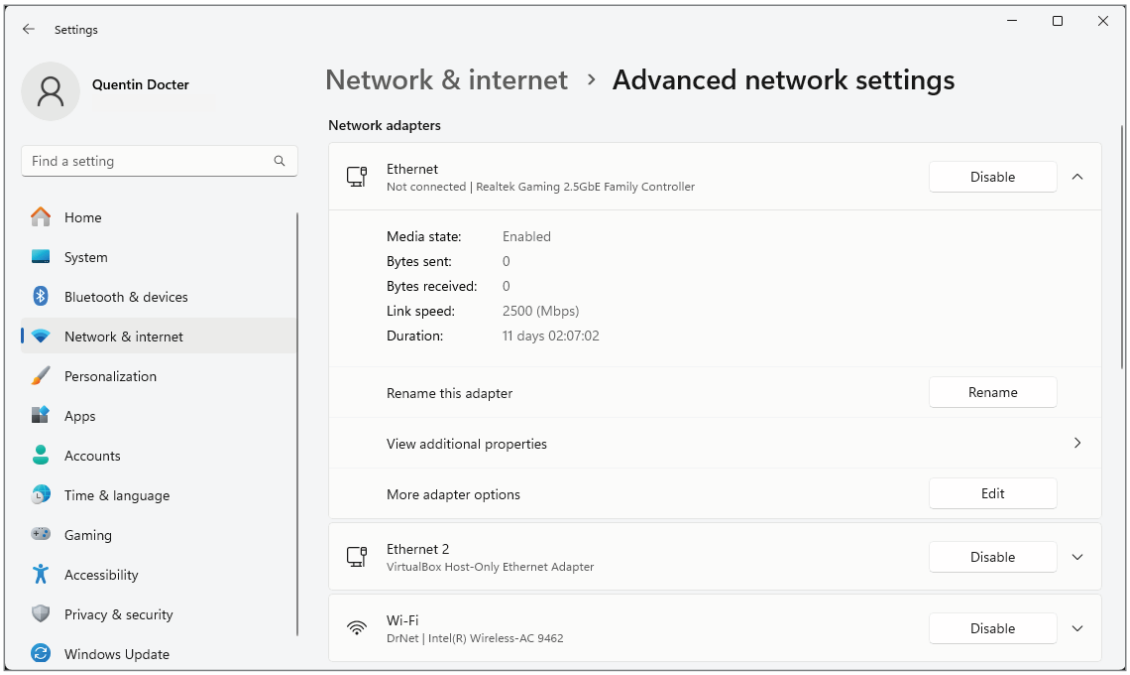

As you can see in Figure 9.12, this computer is configured to obtain its information automatically from a DHCP server. (If you have a wireless router, as many people do on their home networks, it can function as a DHCP server. We'll talk more about that in a few sections.) If you wanted to configure the client manually, you would click Use The Following IP Address and enter the correct information. To supply the client with a DNS server address manually, click Use The Following DNS Server Addresses.

FIGURE 9.12 TCP/IPv6 Properties

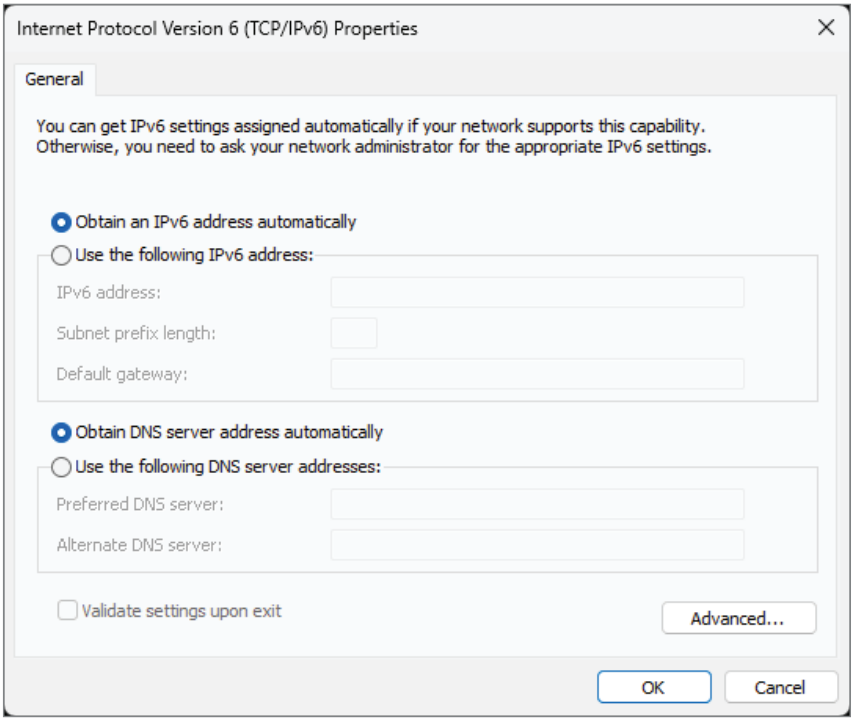
Internet Protocol Version 6 (TCP/IPv6) Properties
General
You can get IPv6 settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IPv6 settings.
Obtain an IPv6 address automatically
Use the following IPv6 address:
IPv6 address:
Subnet prefix length:
Default gateway:
Obtain DNS server address automatically
Use the following DNS server addresses:
Preferred DNS server:
Alternate DNS server:
Validate settings upon exit
Advanced...
OK
Cancel

If you manually configure the IP address, you must also manually configure the DNS server address. Otherwise, the client will not have access to a DNS server and won't be able to use the Internet effectively. Client computers can broadcast to find a DHCP server, but they cannot broadcast to find a DNS server. In addition, don't forget the default gateway, which is the address to the router that takes the network to the outside world. Without the gateway, the computer won't have Internet access!

Configuring a Wireless NIC

Installing an internal NIC is the same whether the card is wired or wireless. Of course, the big difference is how the card connects to the network. A wireless card needs to know the name of the wireless network, or the Security Set Identifier (SSID).To configure a wireless connection, you can simply bring a Windows (XP or newer) laptop or computer within range of a wireless access point, and Windows will automatically detect the access point. Alternatively, if you would like control over the connection, in Windows 11, you can choose Start ➢ Settings ➢ Network & Internet, and then click on Wi-Fi and Show Available Networks (Figure 9.13).

FIGURE 9.13 Available Wi-Fi Networks

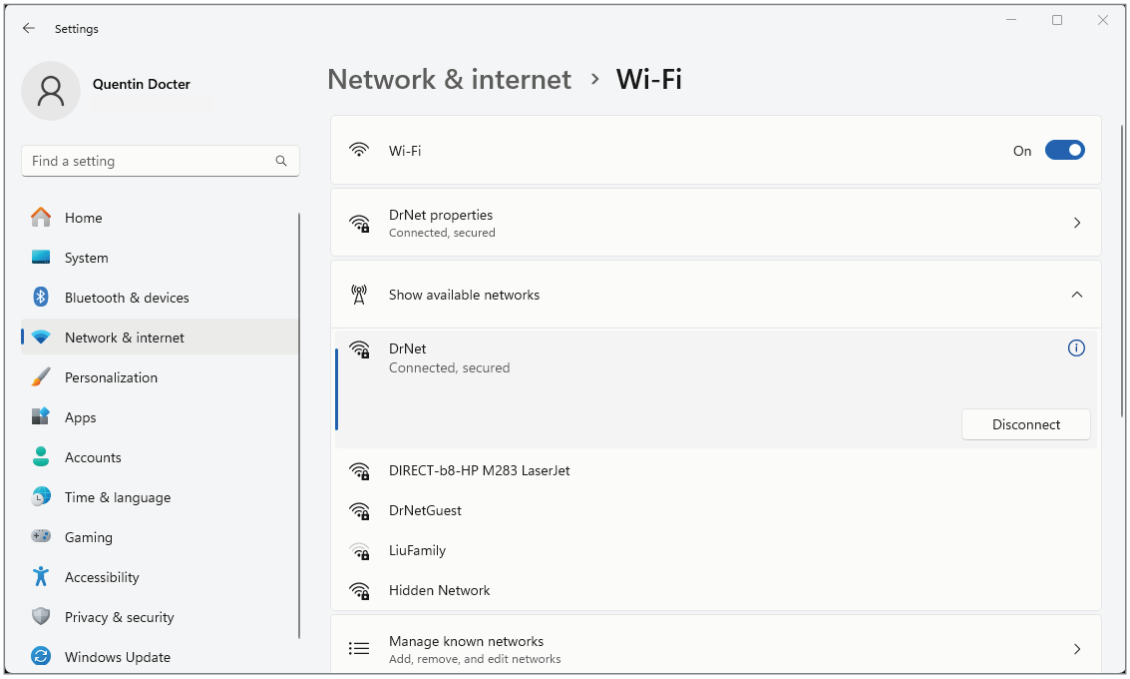

If you have a wireless signal strength indicator in the notification area next to the clock, you can click it to manage available wireless networks as well.

From this screen, you can view the SSIDs of the available wireless networks, including the one to which you are connected (the one that says "Connected" next to it). The icon to the left of the network name indicates the relative signal strength of each connection. Stronger (and faster) connections will have more bars. Also notice that below the list of networks, there is a Manage Known Networks option. This can be used to add a new network name (e.g., if the SSID is hidden), forget existing networks, or configure the computer to not automatically connect to a network that it currently automatically connects to.

If the connection says Secured or shows a lock icon next to the signal indicator (depending on the OS), it is a secured wireless network and you will need to enter some sort of password to gain access to it. It can be dangerous to join unsecured networks; you have no way of knowing who is on them or their intentions. Wireless attacks on unsecured networks are somewhat common, so be careful with joining unsecured networks!

To connect to any network, select it and click the Connect button, and Windows will try to connect. The weaker the signal, the longer the connection will take. Authentication will also slow down the initial connection time. Once you are connected, Windows will display "Connected" next to that connection.

Real World Scenario

Easy Configuration

Imagine that you have a small network of no more than 10 computers and do not have a DHCP server. You want to minimize the administrative hassle of configuring TCP/IP, so you want your computers to configure themselves automatically. What do you do?

The answer is to set the NIC up to get its IP information from the DHCP server anyway. Microsoft Windows operating systems will automatically configure themselves with an Automatic Private IP Addressing (APIPA) address (with the format 169.254.x.x) if they are unable to locate a DHCP server. With an APIPA address, computers on the local network will be able to communicate with one another. The limitation is that the computers will *not* be able to communicate with any remote devices (those not on the local network) and will not be able to get on the Internet.

Installing Network Cables

Network cables are not the most fun things to install. Proper installation of network cables generally means running them through ceilings and walls and making a mess of the office. Thank goodness for wireless!

Be sure to use plenum cable if you are running cables through spaces where there is air ventilation, such as drop ceilings. PVC-coated cables produce poisonous gas when burned. Also, be sure that you have the proper permission to run the cables and that you aren't violating any building codes.

If you are installing a wired network in an existing office space, you may want to look into hiring out the cable installation to a third party. Many companies have the tools to properly install a wired network.

When installing a wired network yourself, always be aware of the maximum cable lengths. In addition, utilize cable troughs in ceilings and walls or another conduit in walls to keep your cables organized. Figure 9.14 illustrates a cable trough; they come in a variety of lengths and quality.

FIGURE 9.14 Cable trough

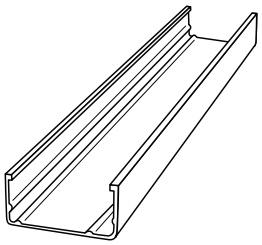

Finally, if you must run cables across the floor in a walkway (which isn't recommended), use a floor cable guard to avoid creating a trip hazard and to protect your cables. A floor cable guard is shown in Figure 9.15.

When running cables through a ceiling, never run the cables directly across fluorescent lights. These lights emit electromagnetic radiation (EMI) that can interfere with network communications. Utilize your cable troughs to keep cables in one place and away from lights. Also remember that fiber-optic cables are immune to EMI.

Also, when installing cables, always be sure to label both ends! That way, you know which cable you're dealing with and where the cable should theoretically go.

FIGURE 9.15 Floor cable guard

Installing and Configuring Connectivity Devices

In this network installation section, we started with the local computer (the NIC) and then moved on to cabling, which, of course, is needed if the network is wired. Continuing our trek away from the local computer, we now need to look at devices that help you connect to other computers. Broadly speaking, we can break these connectivity devices into two categories: those that make connections to the Internet and those that make connections to other local computers. In the first category, we have DSL modems, cable

modems, and ONTs; in the second category, we have switches, wireless routers, and wireless access points.

To be fair, devices that connect computers to other local computers may also get all the computers onto the Internet. For example, your laptop might connect to a wireless router, which lets you talk to a printer on your local network as well as access the Internet.

Installing and Configuring DSL Modems, Cable Modems, and ONTs

We covered the basic installation of these three Internet access devices earlier in the chapter, but since we're talking about network installation, now is a good time to review.

To access the outside world, the Internet access device needs an external connection. DSL modems connect to telephone lines, cable modems to television cable connectors, and ONTs to one or the other with a fiber-optic cable. The ISP manages the connection for you. Internally, the DSL/cable modem or ONT can be connected directly to a computer using a UTP cable, or it can connect to a switch, router, or wireless connectivity device so that multiple users can share the connection.

For the most part, there is little to no configuration you can perform on a DSL/cable modem or ONT. The ISP must initiate the connection to the Internet from their end. Sometimes they need to send a technician to your location to enable it, but other times they just ship the device and you plug it in. Don't forget to plug in the power cord! Beyond that, most ISPs don't want you touching any settings on their devices. If something is wrong, you need to reach out to their tech support people for assistance.

Some DSL and cable modems have built-in wireless router capabilities. If so, it's possible that the ISP will want to charge you more per month for that feature. If the modem is so enabled, you may be able to configure basic settings through a web-based interface. Configuring one of these is very similar to configuring a stand-alone wireless router, which we'll cover in detail later in this chapter.

Installing and Configuring Switches and Hubs

Wired switches and hubs are fairly easy to install. Plug in the power, and plug in the network cables for client computers. Hubs don't typically have configuration options—they just pass signals along. Unmanaged switches are similar to hubs in the sense that there's not much to configure. Managed switches will have configuration settings for VLANs and other options. For more information on the services that managed switches can provide, refer to Volume 1 Chapter 7.

Installing Wireless Access Points and Routers

Instead of using switches and hubs, wireless networks use either a *wireless access point (WAP)* or a *wireless router* to provide central connectivity. A WAP functions essentially like a wireless hub, whereas a wireless router provides more functionality, similar to that of a wired router. Based on looks alone, they are pretty much identical, and physically installing them is similar. The differences come in configuring them because they will have different options.

We're going to talk about installing and configuring WAPs and wireless routers interchangeably; just remember that a lot of the features available in a wireless router may not be available in a WAP.

After unwrapping the device from its packaging (and reading the instructions, of course), you must choose a place for it. If it is supplying wireless access to your home network and the Internet, locate it where you can receive access in the most places. Keep in mind that the more walls the signal has to travel through, the lower the signal strength.

In addition, you may choose to have some computers plug directly into the device using a UTP cable. If so, it makes sense to locate the device near the computer or computers you will want to physically connect.

Place the WAP in the center of your home, close to a network connection. Or if you have only one computer, place it close to the broadband Internet connection you are using (i.e., the fiber-optic, cable modem, or DSL line).

In many offices, WAPs and wireless routers are often placed in the ceiling, with the antennae pointed downward through holes in the ceiling tiles. You can purchase metal plates designed to replace ceiling tiles to hold these devices. The plates have holes precut in them for the antennae to stick through, are designed to securely hold the device and easily open for maintenance, and often lock for physical security. You can also purchase Wi-Fi ceiling antennas that basically look like a little upside-down dome attached to the ceiling.

For wireless connectivity devices placed in a ceiling (or other places with no easy access to an electrical outlet), *Power over Ethernet (PoE)* is a very handy technology to supply both power and an Ethernet connection.

Once you have chosen the location, plug the unit into a wall outlet and connect the antennae that come with the unit (as needed; many newer devices contain built-in antennae). They will screw onto two bungs on the back or side of the unit. Once the unit is plugged in, you need to connect it to the rest of your network.

If you are connecting directly to the Internet through a DSL/cable modem or ONT or to a wired hub or router, you will most likely plug the cable into the Internet socket of the device, provided that it has one. If not, you can use any of the other wired ports on the back of the device to connect to the rest of your network. Make sure that you get a link light on that connection.

At this point, the device is configured for a home network, with a few basic caveats. First, the default SSID (e.g., Linksys) will be used, along with the default administrative password and the default IP addressing scheme. Also, there will be no encryption on the connection. This is known as an *open access point*. Even if you have nothing to protect except for the Internet connection, you shouldn't just leave encryption turned off. It makes you an easy and inviting target for neighbors who want to siphon off your bandwidth or even worse. Many wireless manufacturers have made their devices so easy to configure that for most networks it is Plug and Play.

If you have personal data on your home network and more than one computer, you should never keep the default settings. Anyone could snoop your access point from the road in front of or behind your house and possibly get on your home network. It's too easy for identity theft!

From a computer on the home network, insert the device's setup media (such as a flash drive) into the computer. It will automatically start and present you with a wizard that will walk you through setting the name of the SSID of this new access point as well as changing the default setup password, setting any security keys for this connection, and generally configuring the unit for your network's specific configuration. Then you're done!

Configuring Wireless Routers and Access Points

Each wireless manufacturer uses different software, but you can usually configure their parameters with the built-in, web-based configuration utility that's included with the product. While the software is convenient, you still need to know which options to configure and how those configurations will affect users on your networks. The items that require configuration depend on the choices you make about your wireless network. We will divide the configuration section into two parts: basic configuration and security options, which apply to both routers and access points, and then additional services that are normally router-only.

In the next section, we will refer to wireless routers, but know that the basic configuration options are often available on wireless access points as well.

Basic Configuration

The Wi-Fi Alliance (`www.wi-fi.org`) is the authoritative expert in the field of wireless LANs. It champions five critical steps to setting up a secured wireless router:

1. Change the router's SSID.
2. Change the administrator username and password. Make sure it's a strong password.
3. Select WPA3, WPA2, or AES for security. (We'll discuss these shortly.)
4. Choose a high-quality security passphrase.
5. From the clients, select WPA3 (or WPA2) and enter the security passphrase to connect.

The parameter that needs immediate attention is the SSID. An SSID is a unique name given to the wireless network. All hardware that is to participate on the network must be configured to use the same SSID. Essentially, the SSID is the network name. When you are using Windows to connect to a wireless network, all available wireless networks will be listed by their SSID when you click Show Available Networks.

Setting up a wireless network isn't explicitly called out in the A+ exam objectives. However, given the wireless networking objectives that *are* included, we wouldn't be surprised if you were asked questions related to material covered in this section. For example, SSID is not listed in the 220-1201 objectives, but it *is* in the acronym list at the end of the objectives and also in the 220-1202 objectives. Per CompTIA, test candidates are "encouraged to review the complete list and attain a working knowledge of all listed acronyms." In other words, SSID is probably fair game for the 220-1201 exam.

When you first install the wireless network, the default SSID is used and no security is enabled. In other words, it's easy to find your network (Linksys), and anyone within range of your signal can get on your network with no password required. This is obviously a security risk, so you want to change that.

For the rest of this example, we'll use a Linksys MR9000 wireless router. (Yes, it's only Wi-Fi 5. Sometimes you have to work with what you have. The information in this chapter will still guide you in the right direction on what needs to be done on any Wi-Fi router.) First, you need to log into your device. The default internal address of this router is 192.168.1.1, so to log in, open Chrome, Microsoft Edge, or your preferred Internet browser, and type **192.168.1.1** into the address bar. (Some routers use 192.168.0.1 as a default; check your router's documentation if you are unsure about what your router uses.) You'll get a screen similar to the one shown in Figure 9.16.

Some wireless router installation programs install an icon on your desktop to use for management. Clicking the icon will take you to the management program. If you're using a browser, you can likely manage a wireless router though your phone, too, but the option screens will look a little different.

You should have already set up the username and password using the installation media provided with the device. If not, look in the documentation for the default username and password. You'll definitely want to change these as soon as possible. Once you're logged in, the first screen you'll see is similar to the one shown in Figure 9.17. You can see sections along the left side that allow you to configure various router settings. On this router, the Connectivity section has an Internet Settings tab that identifies how you configure your incoming connection from the ISP. In most cases, your Internet provider will just have you use DHCP to get an external IP address from its DHCP server, but there are options to configure this manually as well.

Next, configure the parameters that are crucial for operation according to the Wi-Fi Alliance. On this router, the admin password is configured on the Basic tab of the Connectivity settings, as shown in Figure 9.18.

FIGURE 9.16 Logging into the wireless router

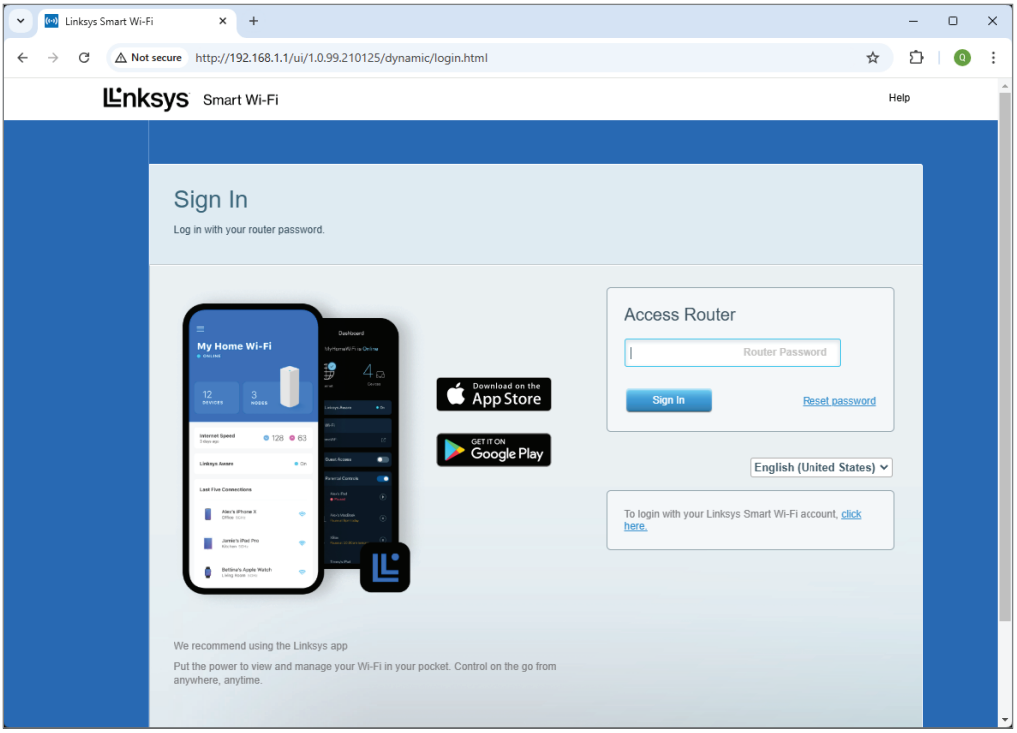
Linksys Smart Wi-Fi
Not secure
http://192.168.1.1/ui/1.0.99.210125/dynamic/login.html
Linksys Smart Wi-Fi
Help
Sign In
Log in with your router password.
Access Router
Router Password
Sign In
Reset password
Download on the App Store
GET IT ON Google Play
English (United States)
To login with your Linksys Smart Wi-Fi account, click here.
We recommend using the Linksys app
Put the power to view and manage your Wi-Fi in your pocket. Control on the go from anywhere, anytime.

FIGURE 9.17 Basic setup screen

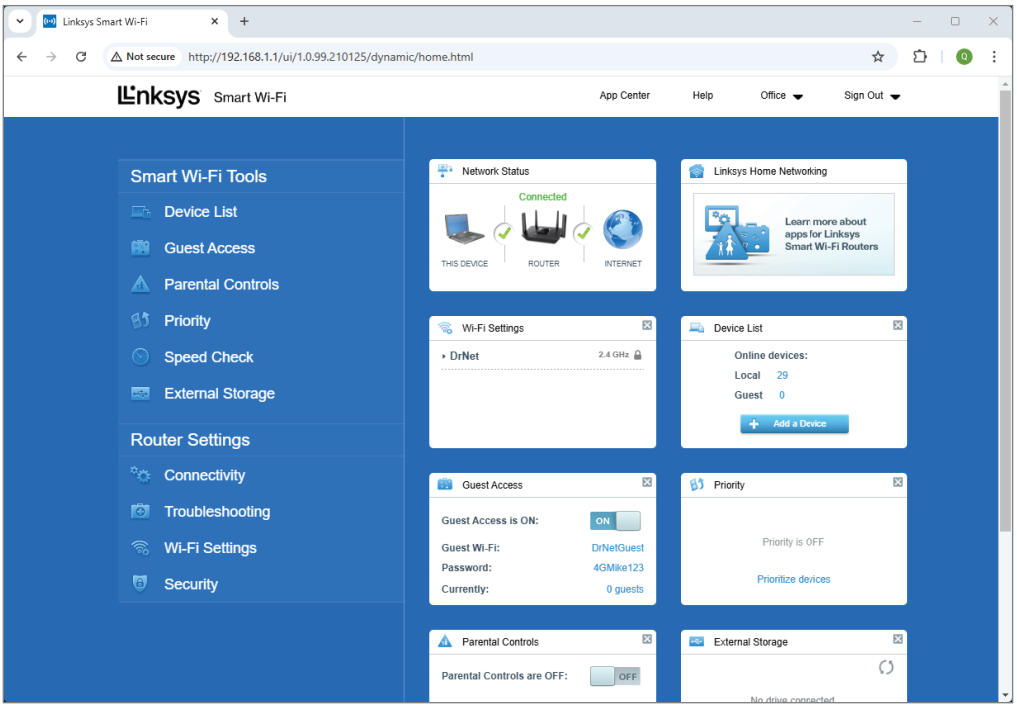

Figure 9.18 also shows the option to check for firmware updates, and in this case those updates will be automatic. After you install a router, change the SSID and passwords, and set up security, you should update the firmware to ensure that it's the most current version. Using older firmware versions could present security risks.

FIGURE 9.18 Basic wireless settings tab

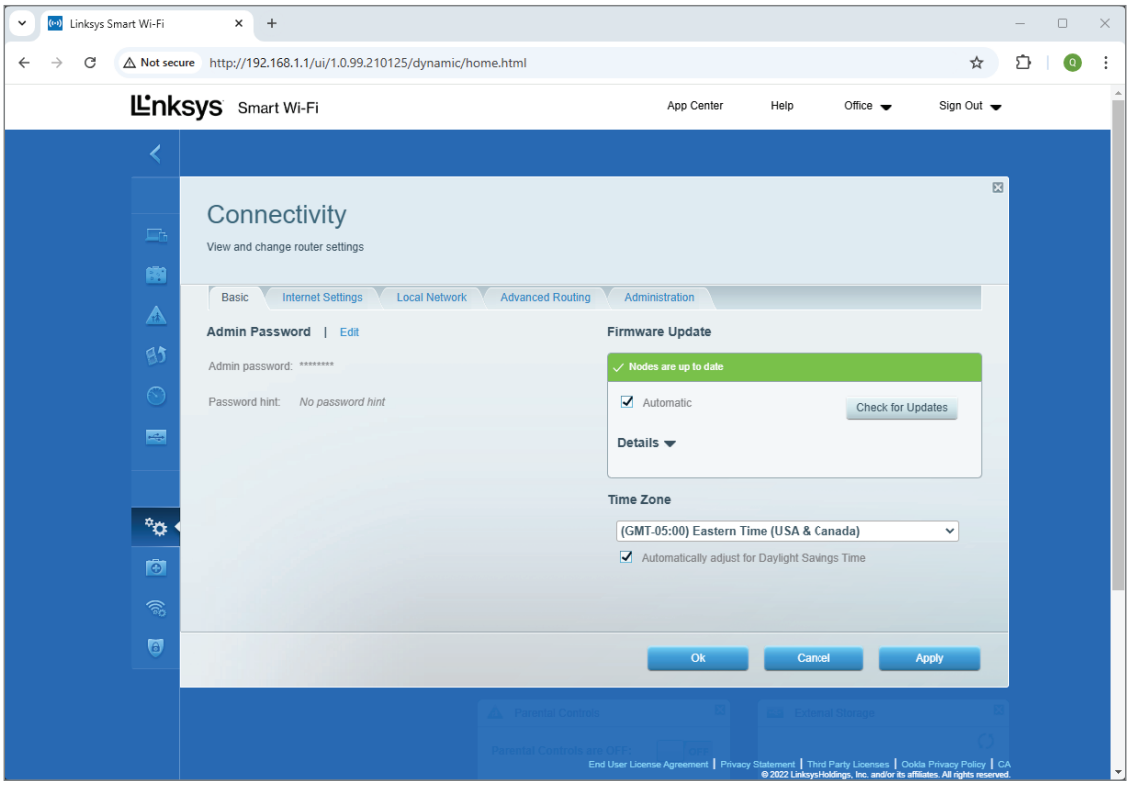

The network name (SSID), the password required by clients to join the network, and the security mode is on the Wi-Fi Settings tab, shown in Figure 9.19. (We blocked out the password for pretty obvious reasons, because this router screen shows it in plain text.) You can change either of these parameters by editing the text in the boxes. Make sure the passwords to join are very different from the administrator password! These steps take care of the SSID, admin password, and security phrase.

Hiding your SSID (by not broadcasting it) does *not* increase network security! It's very easy for someone with a wireless packet sniffer to find an SSID whether it's hidden or not. Hiding the SSID only makes it more difficult for clients to find and connect to the network.

FIGURE 9.19 Wi-Fi Settings tab

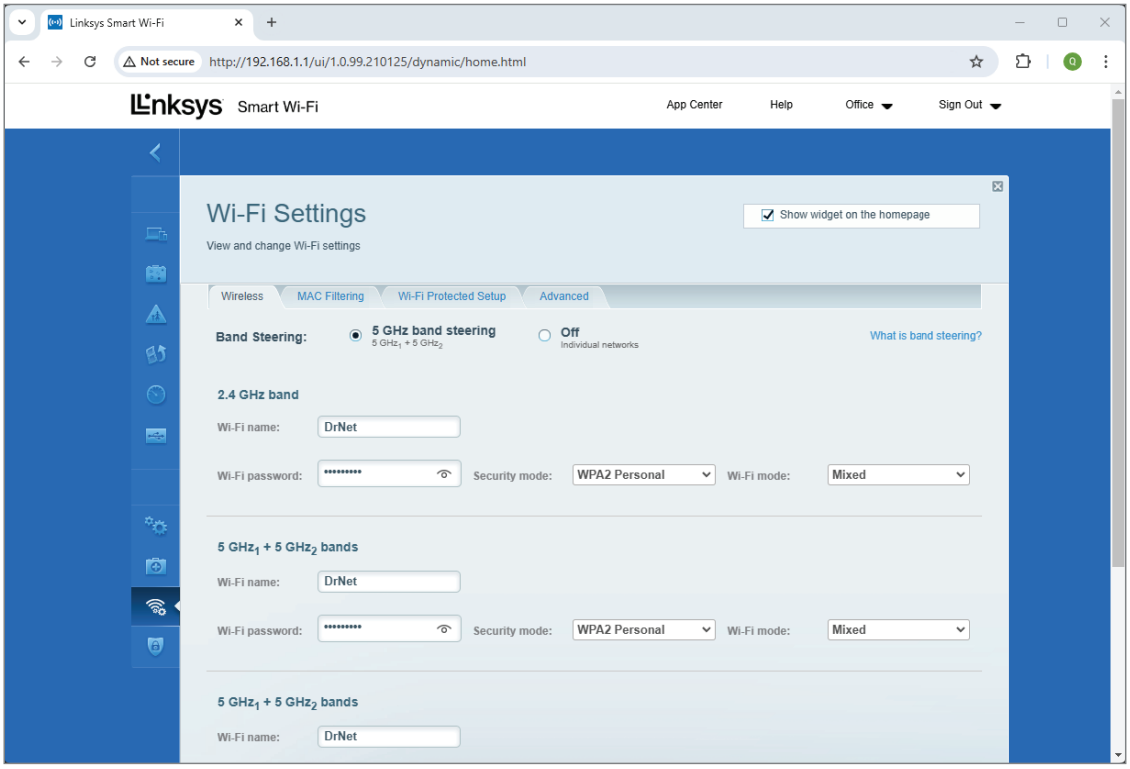

Let's pop back to Connectivity for a minute to configure the internal network settings on the Local Network tab, as shown in Figure 9.20.

Here, you configure your router's hostname, internal IP address (in this case, 192.168.1.1), and subnet mask. On this router, DHCP is also configured on this screen. If you want the device to act as a DHCP server for internal clients, enable it here, specify the starting IP address, and specify the maximum number of DHCP users. Disabling DHCP means that clients will have to use a static IP address.

Most wireless routers (like the one used in this example) have a help section to describe each setting on the configuration pages. If you're not totally sure what a setting does, click the Help link (at the top of the screen in this case) to find out what the setting does. If not, there's always the manual or online help.

The last critical setting you need to make is to enable wireless *encryption*. If you don't do this, all signals sent from the wireless router to client computers will be in plain text and anyone can join the network without a security password. It's a really bad thing to leave disabled. Before we look at how to set up encryption on our wireless router, let's cover the details on three encryption options.

FIGURE 9.20 Local Network Settings screen

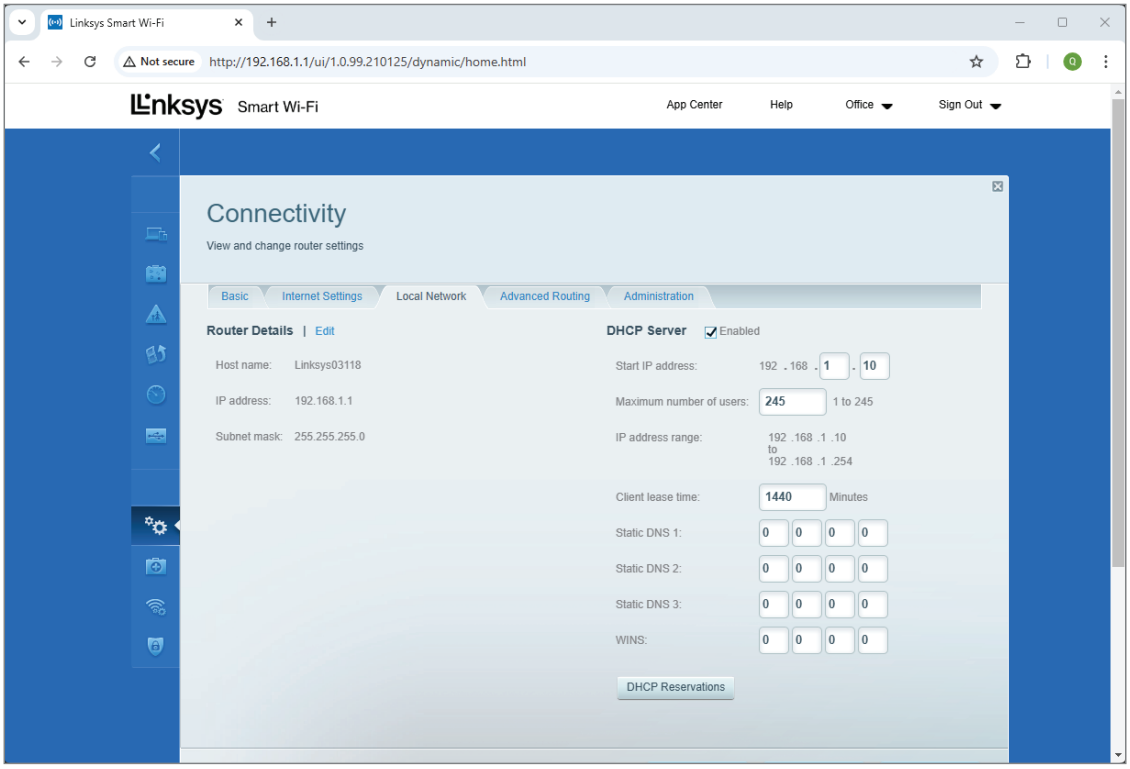

Like the term SSID, the upcoming acronyms WPA, TKIP, AES are part of the A+ acronym list. Be familiar with what they are and what they do.

Wi-Fi Protected Access *Wi-Fi Protected Access (WPA)* was first available in 1999 but did not see widespread acceptance until around 2003. Once it became widely available, the Wi-Fi Alliance recommended that networks no longer use the outdated and flawed Wired Equivalent Privacy (WEP) encryption method in favor of WPA.

This standard was the first to implement some of the features defined in the IEEE 802.11i security specification. Most notably among them was the use of the *Temporal Key Integrity Protocol (TKIP)*. Whereas WEP used a static 40- or 128-bit key, TKIP uses a 128-bit dynamic per-packet key. It generates a new key for each packet sent. WPA also introduced message integrity checking.

When WPA was introduced to the market, it was intended to be a temporary solution to wireless security. The provisions of 802.11i had already been drafted, and a standard that employed all the security recommendations was in development. The upgraded standard would eventually be known as WPA2.

WPA, WPA2, and WPA3 (discussed next) have two variants: Personal and Enterprise. For a small office or home office network with just one wireless router or access point, Personal is the choice to make. With Personal, the device itself handles the authentication. For larger networks, Enterprise is recommended because it consolidates authentication administration. Enterprise requires the use of a separate central authentication server, such as a Remote Authentication Dial-in User Service (RADIUS) server.

Wi-Fi Protected Access 2 Even though their names might make you assume that WPA and WPA2 are very similar, they are quite different in structure. *Wi-Fi Protected Access 2 (WPA2)* is a huge improvement over WPA. As mentioned earlier, it implements all the required elements of the 802.11i security standard. Most notably, it uses Counter Mode CBC-MAC Protocol (CCMP), which is a protocol based on the *Advanced Encryption Standard (AES)* security algorithm. CCMP was created to address the shortcomings of TKIP; consequently, it's much stronger than TKIP.

The terms *CCMP* and *AES* tend to be interchangeable in common parlance. You might also see it written as *AES-CCMP*.

Wi-Fi Protected Access 3 The newest and strongest wireless encryption is *Wi-Fi Protected Access 3 (WPA3)*. It was released in 2018 and is mandatory for all new Wi-Fi–certified devices as of July 2020. It offers more robust authentication and increased cryptographic strength to ensure more secure networks.

WPA3-Personal offers increased security even if users choose passwords that don't meet normal complexity requirements through a technology called Simultaneous Authentication of Equals (SAE). SAE is resistant to dictionary attacks, which are brute-force methods of guessing passwords.

WPA3-Enterprise has a plethora of new security features, including multiple authentication methods, enhanced encryption with 128-bit Advanced Encryption Standard Counter Mode with Cipher Block Chaining Message Authentication (AES-CCMP 128), even better encryption than CCMP in Galois/Counter Mode (AES-GCM), more robust 256-bit security key derivation and confirmation, and 128-bit frame protection. Finally, there's also a 192-bit WPA-3 mode that increases the strength of all features already mentioned. It's way more than you will need to know for the A+ exam, but if cybersecurity is of interest to you, check it out. On this particular router, it would make sense to configure security via the Security section, but that's not true. (Who says things need to be logical?) On this particular router, the encryption method is selected on the same page the SSID and network password are configured. It's in the Wi-Fi settings section, Wireless tab, that was shown in Figure 9.19.

This router happens to be 802.11ac (Wi-Fi 5), so it has sections for both a 2.4 GHz and a 5 GHz network. (Wi-Fi 6E and newer routers will have 6 GHz settings as well.) If there were only devices of one type, it would make sense to disable the other network. In this case, though, we are talking about security, and you can see that it's set to WPA2 Personal. It's always good practice to set up the router to have the most advanced security possible, which is currently WPA3. This router is a bit dated and doesn't support WPA3, so WPA2 is the best choice.

To change the setting, click the down arrow next to Security mode. The other WPA2 choice you generally have is WPA2 Enterprise, which is more secure than Personal. For a business network, regardless of the size, Enterprise is the way to go. In order to use Enterprise, though, you need a separate security server called a RADIUS server.

Another tab in the wireless settings is *MAC Filtering*. On that tab, you can configure the router to allow connections only from clients with certain MAC addresses. (Recall that MAC addresses are hardware addresses built into NICs.) This is a light-duty security measure you can use on your network, though an attacker capable of spoofing a MAC address can easily defeat it.

For your wireless router to use WPA2 Enterprise or WPA3 Enterprise, remember that you need to have a RADIUS server on your network. Remember it this way: WPA Personal means the local (personal) device handles security, whereas WPA Enterprise means some other machine on the network (enterprise) does. The A+ exam won't test you on how to configure RADIUS. For now, just know what it is, which is an authentication server. Other exams, such as Network+, will test you on how to configure RADIUS.

With that, the basic router-side setup recommendations have been taken care of. Now it's just a matter of setting up the clients with the same security method and entering the passphrase.

Real World Scenario

Sharing an Internet Connection

Wireless routers have many advantages over wireless access points. One of the biggest advantages is the ability to share an Internet connection. By sharing a connection, you pay for only one connection, but you can connect as many computers as you would like (or as many as are reasonable) to your wireless router. Here's how to do that.

First, ensure that your Internet connectivity device (DSL/cable modem or ONT) is connected properly. Then, connect your wireless router to that device using a UTP cable (Cat 6 or

better). In most cases, the wireless router will have a wired Internet port on the back. Connect the cable here and plug it into your broadband device. Finally, you can connect computers to your wireless router.

Some ISPs might have the Internet connection tied to a specific MAC address. If you do something like upgrade your wireless router, it might not be able to connect because it will have a different MAC address. You can do one of two things to fix this. The first option is, contact your ISP and have them change the setting. The second option is that most wireless routers will allow you to clone a MAC address (see Figure 9.21). Your new router will simply tell the ISP that it has the same MAC address as your old router, and problem solved.

FIGURE 9.21 Option to clone a MAC address

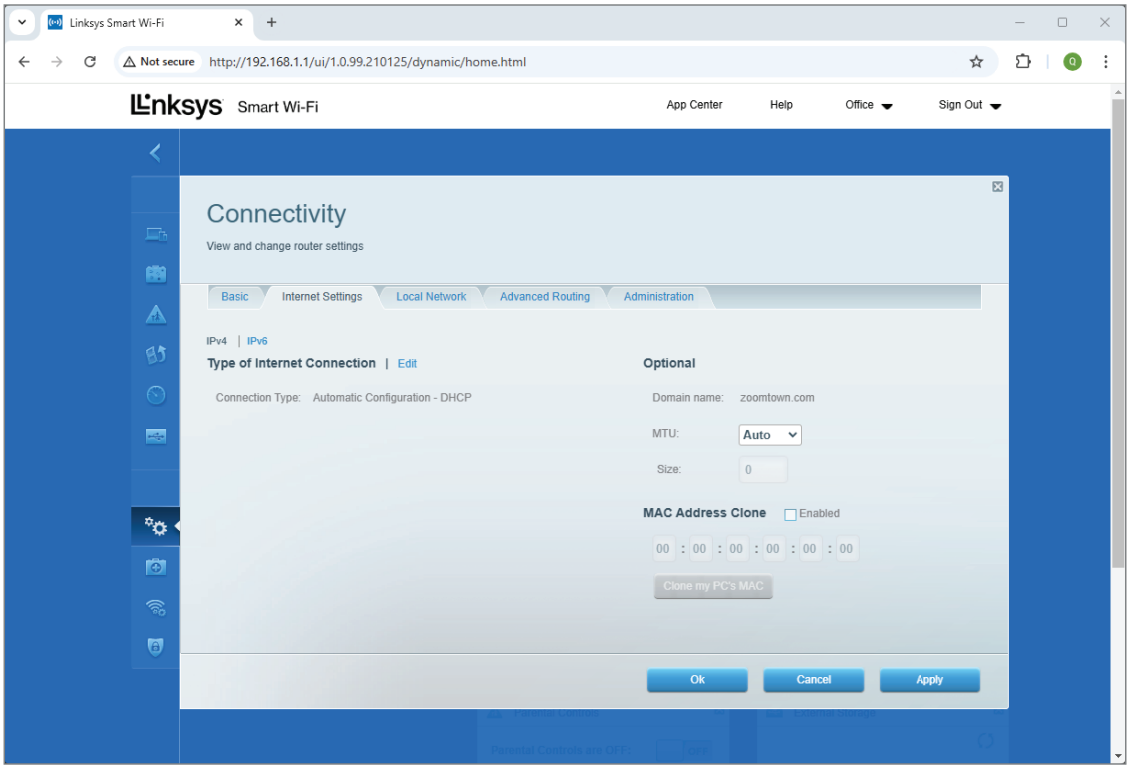

Wireless Channels

Earlier in the chapter, in the section on 802.11g, we brought up the concept of wireless channels, which are divisions within a frequency band. For example, there are 11 configurable channels in the 2.4 GHz range.

In situations where there are too many users for one WAP to adequately service, or the physical layout of the network is too big for one WAP to cover, multiple WAPs are needed. In situations like these, here's how you should configure it:

1. Set up the WAPs so that they have overlapping ranges.

   The minimum overlap is 10 percent, and 20 percent is recommended. This way, if users roam from one area to another, they won't lose their signal.
2. Configure the WAPs with the same SSID.
3. Configure the WAPs with nonoverlapping channels.

Recall that 2.4 GHz channels need to be at least five numbers apart to not overlap. So, for example, channels 2 and 7 do not overlap, nor do channels 4 and 10. There are 11 configurable channels, so you can have a maximum of three overlapping ranges on the same SSID, configured with channels 1, 6, and 11, and not have any interference. Wireless clients are configured to auto-detect a channel by default, but they can be forced to use a specific channel as well.

With Wi-Fi 6 and the introduction of BSS coloring, problems caused by overlapping channels have mostly become a thing of the past. And most routers will automatically configure channels to avoid conflicts anyway. Still, you should know how to change the channel if needed. For example, Figure 9.22 shows a wireless router configuration page where you can configure the 5 GHz network. In this case, you can choose from 20 MHz or 40 MHz channel widths, as well as choose the channel. Each of the 20 MHz channels shown is nonoverlapping.

FIGURE 9.22 5 GHz channels available to select

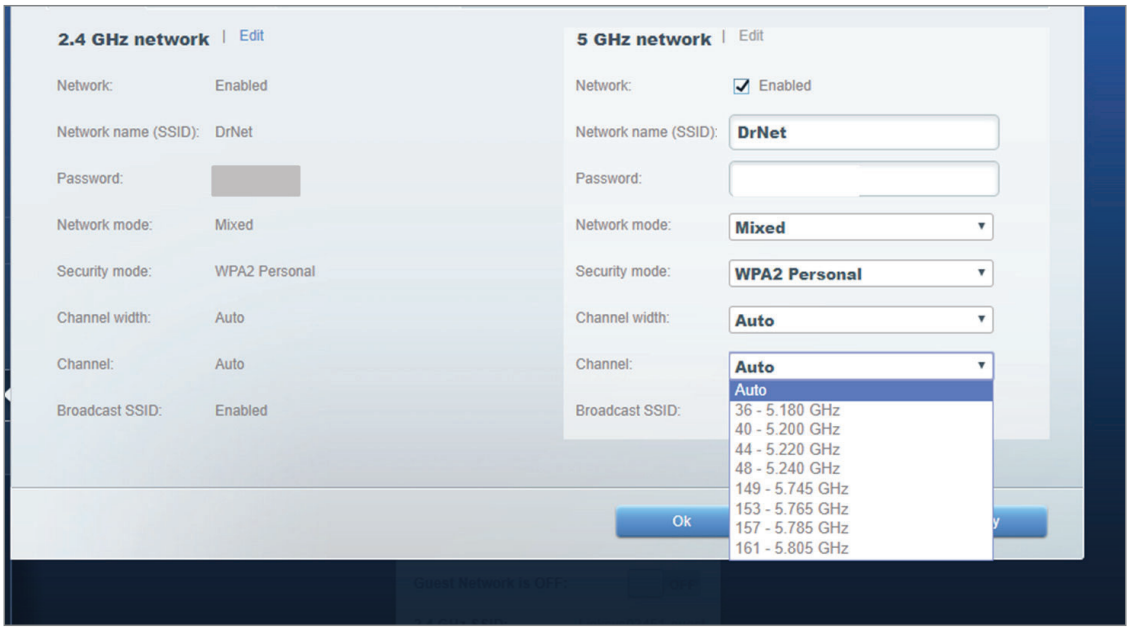

Additional Wireless Router Services

Wireless routers offer many more services than access points offer. Most of those services fall under the umbrella of configuring your router as a firewall, which we briefly introduced in

Volume 1 Chapter 6. Firewalls are an important networking concept to understand, so let's dive deeper into what firewalls are and what they do, and then we'll look at specific configuration options.

Understanding Firewall Basics

Before we get into configuring your wireless router as a firewall, let's be sure you know what firewalls can do for you. A *firewall* is a hardware or software solution that serves as your network's security guard. For networks that are connected to the Internet, firewalls are probably the most important device on the network. Firewalls can protect you in two ways. They protect your network resources from hackers lurking in the dark corners of the Internet, and they can simultaneously prevent computers on your network from accessing undesirable content on the Internet. At a basic level, firewalls filter packets based on rules defined by the network administrator.

Firewalls can be stand-alone "black boxes," software installed on a server or router, or some combination of hardware and software. Most firewalls have at least two network connections: one to the Internet, or *public side*, and one to the internal network, or *private side*. Some firewalls have a third network port for a second semi-internal network. This port is used to connect servers that can be considered both public and private, such as web and email servers. This intermediary network is known as a *screened subnet*, an example of which is shown in Figure 9.23. Personal software-based firewalls will run on computers with only one NIC.

FIGURE 9.23 A network with a screened subnet

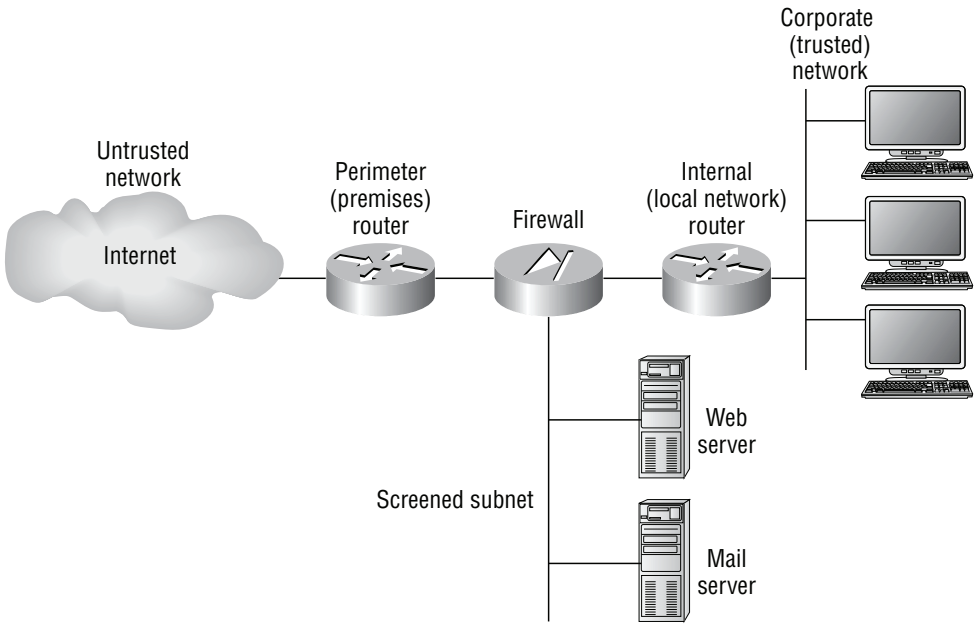

Types of Firewalls

We've already stated that firewalls can be software- or hardware-based or a combination of both. Keeping that in mind, there are two general categories of firewalls: network-based and host-based.

Network-Based Firewalls A *network-based firewall* is what companies use to protect their private network from public networks. The defining characteristic of this type of firewall is that it's designed to protect an entire network of computers instead of just one system. It's generally a stand-alone hardware device with specialized software installed on it to protect your network.

Host-Based Firewalls In contrast to network-based firewalls, a *host-based firewall* is implemented on a single machine, so it protects only that one machine. This type of firewall is usually a software implementation because you don't need any additional hardware in your personal computer to run it. All current Windows client operating systems come with Windows Defender Firewall (or just Windows Firewall), which is a great example of a host-based solution. Norton Security and many other security products come with software firewalls too. Host-based firewalls are generally not as secure as network firewalls, but for small businesses or home use, they're an adequate, cheap solution.

How Firewalls Work

Firewalls are configured to allow only packets that pass specific security restrictions to get through them. They can also permit, deny, encrypt, decrypt, and proxy all traffic that flows through them, most commonly between the public and private parts of a network. The network administrator decides on and sets up the rules a firewall follows when deciding to forward data packets or reject them.

The default configuration of a firewall is generally *default deny*, which means that all traffic is blocked unless specifically authorized by the administrator. While this is very secure, it's also time consuming to configure the device to allow legitimate traffic to flow through it. The other option is *default allow*, which means all traffic is allowed through unless the administrator denies it. If you have a default allow firewall and don't configure it, you might as well not have a firewall at all.

You will also encounter the terms whitelist and blacklist in relation to network and firewall security. A whitelist (or *allow list*) allows access to nobody (so its default is deny) except those users or computers on the whitelist. A blacklist (or *block list*) allows access to everyone (default allow) except those on the blacklist.

The basic method of configuring firewalls is to use an *access control list (ACL)*. The ACL is the set of rules that determines which traffic gets through the firewall and which traffic is blocked. ACLs are typically configured to block traffic by IP address, port number, domain name, or some combination of all three.

Packets that meet the criteria in the ACL are passed through the firewall to their destination. For example, let's say you have a computer on your internal network that is set up as a web server. To allow Internet clients to access the system, you need to allow data on port 80 (HTTP) and 443 (HTTPS) to get to that computer.

Another concept you need to understand is *port triggering*. It allows traffic to enter the network on a specific port after a computer makes an outbound request on that specific port. For example, if a computer on your internal network makes an outbound RDP request (port 3389), subsequent inbound traffic destined for the originating computer on port 3389 would be allowed through.

Configuring Your Wireless Firewall

Nearly every wireless router sold today provides you with some level of firewall protection. On the router used in this example, the firewall options are on two separate tabs. Enabling the firewall and setting a few basic options is done from the Security section, as shown in Figure 9.24. More advanced options, such as configuring port forwarding and port triggering, are on the DMZ and Apps and Gaming tabs. (Remember that DMZs are the old name for screened subnets.)

FIGURE 9.24 Enabling the firewall

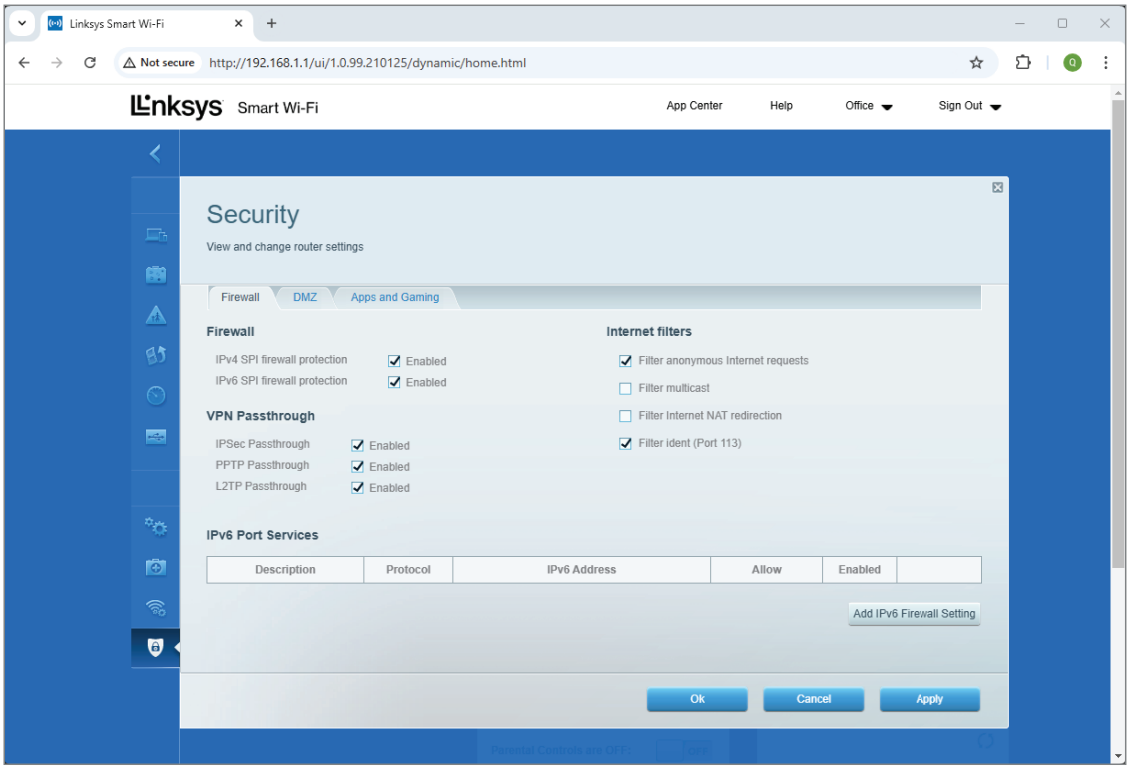

Network Address Translation

Network Address Translation (NAT) is a very cool service that translates a private IP address on your internal network to a public IP address on the Internet. If you are using your wireless router to allow one or more clients to access the Internet but you have only one external public IP address, your router is using NAT.

Most routers have NAT enabled by default, and there might not be any specific configuration options for it. That's true in the case of the router we've been using as an example. You can enable or disable it on the Advanced Routing tab in Connectivity, but otherwise the only options you can configure are the internal IP addresses that the router hands out to clients.

To be technically correct, NAT is specifically a one-to-one translation of a private IP address to a public IP address. If you have multiple client computers with private addresses accessing the Internet using one public address (called many-to-one), that is a specific form of NAT known as overloading, Port Address Translation (PAT), or *port forwarding*. The A+ exam does not test you on the differences between NAT and PAT, but other tests do, such as the Network+ exam.

Another type of NAT is called *Dynamic Network Address Translation (DNAT)*, which translates a group of private addresses to a pool of routable addresses. This is used to make a resource that's on a private network available for consumption on public networks by appearing to give it a publicly available address. For example, if a web server were behind a NAT-enabled router and did not have its own public IP address, it would be inaccessible to the Internet. DNAT can make it accessible.

Universal Plug and Play

Universal Plug and Play (UPnP) is a standard designed to simplify the process of connecting devices to a network and enable those devices to automatically announce their presence to other devices on the network. If you remember when Plug and Play was new to computers, it was revolutionary. You simply plugged in a peripheral (such as a USB network card or mouse) and it was detected automatically and it worked. UPnP is the same idea, but for networking. From a standards standpoint, there's not a lot to configure. The client needs to be a DHCP client and the service uses UDP port 1900.

The concept is great. It lets devices connect to the network and discover each other automatically with the Simple Service Discovery Protocol. It can be used for any networked device you can think of, from routers and printers to smartphones and security cameras.

The problem is, UPnP has no authentication mechanism. Any device or user is trusted and can join the network. That is obviously a problem. The security consulting firm Rapid7 did a six-month research study in early 2013 and found that over 6,900 network-aware products, made by 1,500 different companies, responded to public UPnP requests. In total, they found nearly 81 million individual devices responded to requests. The U.S. Department of Homeland Security and many others immediately began requesting people to disable UPnP.

Since that time, the Open Connectivity Foundation (OCF; `www.openconnectivity.org`), who designs UPnP, has worked to patch security holes. This is a good real-life example of why you should always update the firmware on routers!

The biggest risk is for open UPnP connections to be exploited by unknown systems on the Internet. Therefore, you should configure your router to not allow UPnP connections from its external connection. Many ISPs have also taken steps to help prevent issues. In summary, the best bet is to leave it disabled.

Real World Scenario

Knowing How to Install and Configure a SOHO Router

If you are given the scenario in real life or on the A+ exam, you should know how to install and configure a SOHO router. In today's environment, most installations for small offices and home networks will involve wireless routers. Throughout this chapter, you have learned everything you need to know to purchase the best device for different situations and how to set it up once you get it out of the box. Let's review here.

The first decision to make is which router technology to purchase. With wireless, it's generally best to go with the best technology available because speeds have been increasing rapidly over the last several years. Go with Wi-Fi 7 if possible, with Wi-Fi 6 being an acceptable fallback option.

Next, how will the office get on the Internet? The most obvious choices are fiber-optic, DSL and cable Internet. Find the best solution (which often means the least expensive), and set it up through the service provider. In fact, you may want to do this first, because it will probably take the provider at least a few days to get it set up on their end, and you want it to be ready when you install the network.

Then, as needed, plan the internal network layout. Is it all wireless, or will there be wired connections? Do you need to run cables? Will there be a server closet or other central connectivity point? Once you're certain of your configuration, you can begin installation.

When installing the router, always remember these key steps:

1. Change the default SSID.
2. Change the administrator username and password to something highly secure.
3. Configure the best security possible, such as WPA3.
4. Set a strong passphrase for clients to join the network.

After completing these steps, you can configure the clients to join the network by setting their security appropriately, finding the SSID, and entering the passphrase. By following these guidelines, you will be able to properly install a network, regardless of the scenario presented.

Summary

In this chapter, you learned about wireless networking and configuring a small office, home office (SOHO) network. We started with wireless networking. We introduced the key wireless networking standards all the way through Wi-Fi 7 (802.11be), and talked about their characteristics, such as speed, distances, frequencies, channels, and modulation. Then we moved on to Bluetooth networking, and RFID and NFC.

Next, you learned the fundamentals of installing a small network. We started by looking at network planning, which is the critical first step. Don't be the one who forgets that!

Then, we covered the myriad possibilities for Internet connections, starting with wired broadband options such as fiber-optic, DSL, and cable modems. Then we looked at wireless choices: satellite, cellular, and a wireless Internet service provider (WISP). Then we talked about choosing internal network connections in both wired and wireless environments.

From there, we dove into installing network infrastructure. If you did a good job planning, this part should be problem-free. We covered installing and configuring NICs (including IP addressing for clients), cables, and connectivity devices.

Finally, we looked at how to configure a router. The Wi-Fi Alliance has some great practical steps on how to configure a secure wireless network, such as changing the SSID, setting passwords, and enabling encryption, such as WPA2 and WPA3. We also looked at other basic configuration options, such as DHCP and communication channels. Then we looked at your wireless router as a firewall, including NAT and UPnP.

Exam Essentials

Know the different 802.11 standards. Standards you should be familiar with are 802.11a, 802.11b, 802.11g, 802.11n, 802.11ac (Wi-Fi 5), 802.11ax (Wi-Fi 6), and 802.11be (Wi-Fi 7). Know the frequencies (2.4 GHz, 5 GHz, and 6 GHz) that each one uses, as well as performance characteristics of each, such as relative distance and speed.

Know the three nonoverlapping 2.4 GHz wireless channels. If you need three nonoverlapping channels, you must use channels 1, 6, and 11.

Understand Bluetooth networking. Bluetooth isn't used for wireless LAN like 802.11, but it is used for small personal area networks. It's best for peripheral connectivity such as headsets, keyboards, mice, and printers.

Understand the difference between RFID and NFC. Both use radio signals. RFID is used to track the presence or location of items. NFC uses high-frequency RFID signals and can be used for touchless payment systems.

Know the different types of available broadband connections. Broadband connections include DSL, cable, fiber, satellite, cellular, and wireless Internet service provider (WISP).

Know the various cellular networking standards. Understand the differences between 4G and 5G, and also how 5G has LTE and mmWave.

Know how to configure a network client to use IP addressing. Clients can be configured with static IP information, or dynamically through a DHCP server. To communicate on the Internet, clients need an IP address, subnet mask, and gateway. On an internal network, you can use private IP addresses as long as they go through NAT to get to the Internet. APIPA is available as a fallback if a DHCP server is not available.

Understand the three encryption protocols used for wireless networking. Listed in order from least to most secure, the common wireless security protocols include WPA, WPA2, and WPA3. WPA uses TKIP, WPA2 uses AES, and WPA3 uses AES-GCM.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. Which of the following legacy wireless IEEE standards operate on the 2.4 GHz radio frequency and are directly compatible with each other? (Choose two.)
   - **A.** 802.11a
   - **B.** 802.11b
   - **C.** 802.11ac
   - **D.** 802.11g

2. What is the primary function of the SSID?
   - **A.** To secure communication between a web server and a browser
   - **B.** To secure communication between a server and a remote host
   - **C.** To serve as a parameter used to identify a network and configure a wireless connection
   - **D.** To serve as a type of password used to secure a wireless connection

3. Which of the following are features that allow 802.11ac to achieve higher data throughput? (Choose two.)
   - **A.** MIMO
   - **B.** Beamforming
   - **C.** Channel bonding
   - **D.** Code division multiplexing

4. What of the following is the most secure wireless encryption standard for 802.11 networks?
   - **A.** WEP
   - **B.** WPA
   - **C.** WPA2
   - **D.** WPA3

5. You are upgrading a wireless network from Wi-Fi 5 over to Wi-Fi 6. Which of the following statements is *not* correct?
   - **A.** Users should get faster speeds.
   - **B.** Network range will be similar.
   - **C.** Fewer users will be able to use each access point.
   - **D.** Devices may have lower power usage while using the network.

6. You have just installed a wireless router on your home network. Which of the following should you do to make it highly secure? (Choose all that apply.)
   - **A.** Change the default administrator name and password.
   - **B.** Change the SSID.
   - **C.** Enable WEP.
   - **D.** Configure it to channel 11.

7. You are setting up a small office network for a client. Which Internet service would you recommend to provide the best speed?
   - **A.** DSL
   - **B.** Cable
   - **C.** Satellite
   - **D.** Fiber

8. Which service allows users with private IP addresses to access the Internet using a public IP address?
   - **A.** DHCP
   - **B.** DNS
   - **C.** DSL
   - **D.** NAT

9. You are configuring a single 802.11g wireless network that has been having intermittent communications issues. The office space is large enough that it has three WAPs. What channels should you configure the WAPs on to avoid communication issues?
   - **A.** 2, 5, and 7
   - **B.** 1, 8, and 14
   - **C.** 1, 6, and 11
   - **D.** 3, 6, and 9

10. You are setting up a wireless network. Which wireless standards would give the users over 1 Gbps throughput? (Choose all that apply.)
    - **A.** 802.11ac
    - **B.** 802.11be
    - **C.** 802.11g
    - **D.** 802.11n

11. You have been asked to configure a network for a small office. The wireless router is installed, and now you need to connect the client computers. What do you enter on the client computers to connect to the router?
    - **A.** The administrator password
    - **B.** The security passphrase
    - **C.** The client's MAC address
    - **D.** The default router password

12. Which of the following technologies can operate in the 125 kHz to 134 kHz range?
    - **A.** Bluetooth
    - **B.** RFID
    - **C.** NFC
    - **D.** LTE

13. Due to channel interference, you are upgrading a wireless network from Wi-Fi 5 to Wi-Fi 6. Which feature of Wi-Fi 6 reduces channel interference?

A. OFDMA

B. BSS coloring

C. MU-MIMO

D. DFS

14. Which of the following security standards was the first to introduce a dynamic 128-bit per-packet security key?

A. WEP

B. TKIP

C. AES

D. CCMP

15. Which was the first Wi-Fi version to support WPA3 and OFDMA?

A. Wi-Fi 5

B. Wi-Fi 6

C. Wi-Fi 6E

D. Wi-Fi 7

16. When enabled, which feature of a wireless router allows only specified computers to access the network?

A. Port forwarding

B. NAT

C. SSID

D. MAC filtering

17. A firewall operates by using a set of rules known as what?

A. SLA

B. ACL

C. NAT

D. APIPA

18. You have set up a wireless router on your network and configured it to use AES. What configuration option do you need to choose on the client computers?

A. WEP

B. WPA

C. WPA3

D. TKIP

19. Besides 802.11 standards, which wireless communication method also works in the 2.4 GHz range?

A. Bluetooth

B. Satellite

C. mmWave

D. NFC

20. You are configuring a SOHO client to use TCP/IP. Which parameter is needed to tell the client where to communicate to get on the Internet?

A. Static address

B. Dynamic address

C. APIPA

D. Gateway

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exams require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answers compare to the authors', refer to Appendix B.

You are installing a SOHO network for a client who has purchased copper-based Internet from a major media provider. The hardware shown in the picture is available to you. How do you set it up so that multiple devices can access the Internet? (You may not need to use all of the available hardware.)

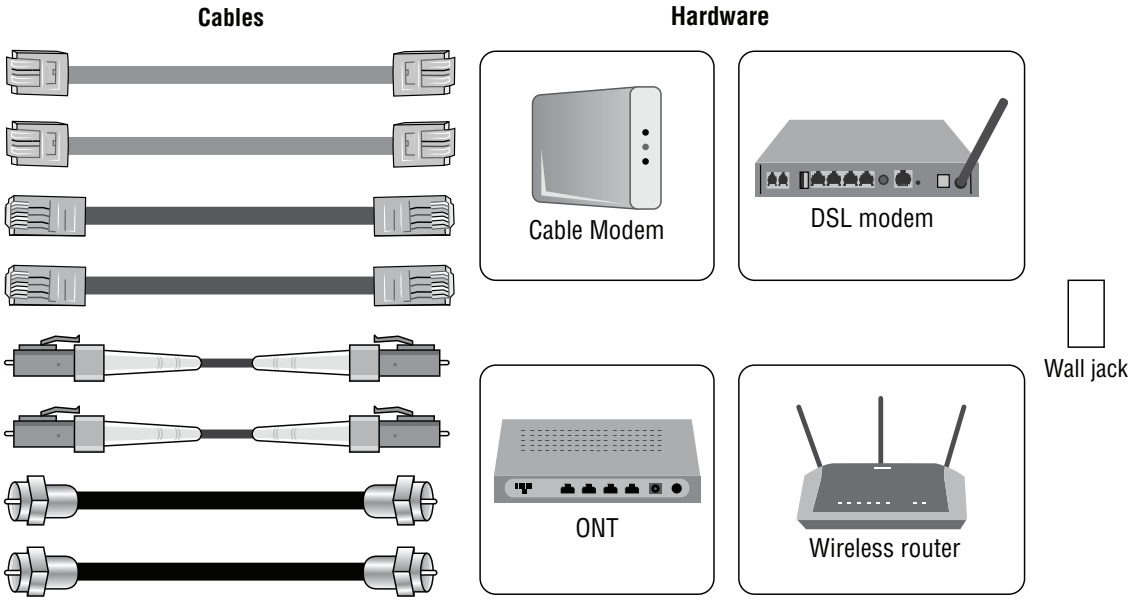

Chapter 10

Mobile Connectivity and Application Support

THE FOLLOWING COMPTIA A+ 220-1201 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

✓ **1.0 Mobile Devices**

✓ **1.3 Given a scenario, configure basic mobile device network connectivity and provide application support.**

- Wireless/cellular data network (enable/disable)
  - 3G/4G/5G
  - Hotspot
  - Wi-Fi
  - Subscriber Identity Module (SIM)/eSIM
- Bluetooth
  - Enable Bluetooth
  - Enable pairing
  - Find a device for pairing
  - Enter the appropriate personal identification number (PIN) code
  - Test connectivity
- Location services
  - Global Positioning System (GPS) services
  - Cellular location services
- Mobile device management (MDM)
  - Device configurations
    - Corporate
    - Bring your own device (BYOD)

- Policy enforcement
- Corporate applications

- Mobile device synchronization
  - Recognizing data caps
  - Calendar
  - Contacts
  - Business applications
    - Mail
    - Cloud storage

Devices that fit into the palm of your hand today are substantially more powerful than the bulky desktop systems of the 1990s, and those were exponentially more powerful than the room-sized supercomputers of the 1960s. This movement toward smaller devices creates new needs, specifically in the area of interaction with the devices—and that relates directly to the operating system and hardware. Remember that the OS is the interface between the hardware and the user, meaning that the OS needs to interpret user input and translate that into an action for the underlying hardware. We won't cover the details of mobile OSs in this chapter—that's taken care of in Volume 2 Chapter 1, "Operating System Basics." For this chapter, you just need to remember that Android and iOS are the two dominant mobile operating systems in the market. Odds are you are familiar with using at least one of them. We'll discuss connectivity and synchronization options for each.

In addition, small mobile devices just don't have room for traditional desktop or laptop hardware. Smartphones don't come with built-in keyboards and mice, nor do small devices have the storage capacity of their larger cousins. Yet with all of their physical space constraints, mobile devices are still expected to perform many of the same tasks as laptops or desktop computers, particularly network connectivity and data storage.

You already know that mobile devices can get on a network and store data. They just do it in a different way than larger systems do. In addition, because of their limited storage capacities, mobile devices can greatly benefit from the syncing of data to a larger storage system, such as a desktop, the cloud, or even an automobile. This chapter focuses on the details of mobile device connectivity, with a particular emphasis on email, as well as device synchronization.

Understanding Mobile Connectivity

Mobile devices like smartphones, tablets, and wearable devices basically exist for their convenience and ability to connect to a network. No one is using their small mobile devices to crunch spreadsheets, create detailed presentations, or perform other tasks that are better suited for larger PCs and Macs. On the other hand, mobile devices are great for surfing the Internet, messaging friends, taking pictures and sending them to friends and family, listening to music and watching videos, and even monitoring our health. Each of these tasks requires a network connection, either wired or wireless, at some point.

Of course, mobile devices are well known for their cellular connectivity, including the ability to access data services over the cellular network. Many subscribers pay a premium for data access, and going over your limit on a data plan can be a very expensive mistake (unless

your plan offers unlimited data). To ease the expense involved in data-network access, nearly all manufacturers provide alternate access methods in their devices. For example, the service provider levies no additional expense for Wi-Fi or Bluetooth data access.

There are other cellular options on mobile devices with which you may be less familiar, though, such as hotspots and tethering and updating the firmware. In the following sections, we'll dive into those topics in some detail.

Connecting to a Wi-Fi network gives the mobile device access to resources on the network, such as printers, with the added bonuses of Internet access (assuming the wireless network has it), free texting, and perhaps Wi-Fi phone calls. Not only will this likely give the mobile user faster speeds than cellular, it also has the benefit of not counting against the data plan.

Bluetooth access to other devices is designed mainly for short-range communications, such as between the mobile device and a nearby computer with which it can exchange or synchronize data. Other applications of Bluetooth generally do not involve the exchange of data but rather the use of a nearby resource such as a printer to make a hard copy of a document or a headset to make hands-free phone calls.

As you can see, there is a distinct advantage to being able to connect your mobile device to a noncellular network. Mobile devices will use Wi-Fi when connected instead of cellular for data operations. Android devices will let you choose which apps can use Wi-Fi only or if cellular is okay to use. Figure 10.1, for example, shows the Google Play Store app settings on an Android device that asks if the app can use Mobile data or Wi-Fi, Wi-Fi only, or Mobile data only.

FIGURE 10.1 Google Play Store app settings

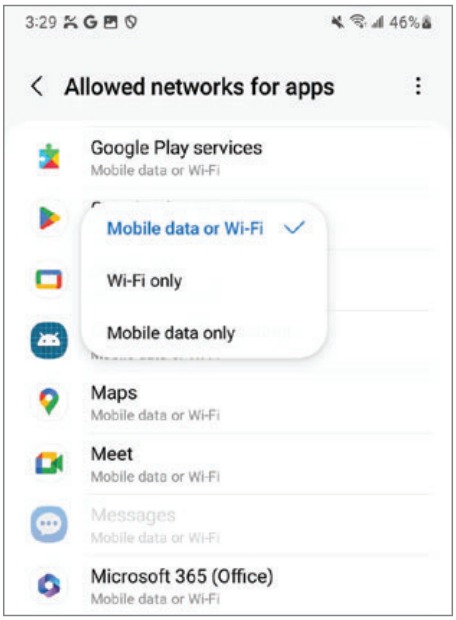

The following sections detail concepts relating to cellular networking and attaching to noncellular networks on iPhones and Android devices. After that, you will be introduced to

the tasks required to establish email connectivity over these mobile units to corporate and ISP connections as well as integrated commercial providers.

Understanding Cellular Networking Standards

Most of the networking technologies we cover in this book are relatively short range. Wi-Fi and Bluetooth signals can stretch a few hundred meters at best, and NFC is even more limited than that. Cellular is by far the longest-range of the wireless networking technologies available today.

We introduced fourth generation (4G) and fifth generation (5G) in Volume 1 Chapter 9, "Wireless and SOHO Networking." We'll review those two here, and also introduce a few of their predecessors.

Third Generation (3G)

Two major cell standards dominated the industry for the formative years of the mobile phone evolution. The *Global System for Mobile Communications (GSM)* was the most popular, boasting over 1.5 billion users in 210 countries. The other standard was *code-division multiple access (CDMA)*, which was developed by Qualcomm and available only in the United States.

3G, 4G, and 5G are network types you need to be familiar with as part of Exam Objective 1.3, "Given a scenario, configure basic mobile device network connectivity and provide application support."

Both were *third generation (3G)* mobile technologies, and each had its advantages. GSM was introduced first, and when CDMA was launched, it was much faster than GSM. GSM eventually caught up, though, and the two ended up with relatively similar data rates. The biggest issue was that GSM and CDMA were not compatible with each other. Whatever technology you got was based on the provider you signed up with. Sprint (which merged with T-Mobile in 2020) and Verizon used CDMA, and AT&T and T-Mobile used GSM. That meant that if you had a CDMA phone through Verizon, you couldn't switch (with that phone) to AT&T. And your CDMA phone wouldn't work outside the United States.

When 3G was first introduced in 1998, it specified a minimum data download rate of 200 Kbps. Data rates increased over time, and some carriers claimed to deliver over 7 Mbps downloads—although technically that was with 3.5G. Really, data rates varied by carrier, the equipment installed in the area, the number of users connected to the tower, and the distance from the tower.

Because of relatively slow data rates (by today's standards), 3G was good for voice calls and texting, but not so great for Internet use.

Prior to 3G there was a 2G digital standard and a 1G analog standard as well. 1G was launched in 1979, and 2G became commercially available in 1991. New standards appear about every 10 years. After several enhancements, 2G had a maximum data rate of 64 Kbps and with it you could send text messages, read emails, and in theory access the Internet. It was very slow and had limited network range, which is why 3G was greatly appreciated when it launched.

Fourth Generation (4G)

In 2008 *fourth generation (4G)* came into the market. Again, we already introduced 4G in Volume 1 Chapter 9 so we won't go into too much depth here, but simply review the key points:

- 4G is designed to use IP instead of traditional telephone circuits, like 3G did. It's designed to provide mobile broadband access.
- In the early days of 4G, there were two competing standards—WiMAX and Long-Term Evolution (LTE). WiMAX initially had higher speeds, but cell phone providers got behind LTE, which won out.
- The biggest enhancement of 4G LTE over 3G is speed. Some 4G LTE networks claim to give you download speeds of 10–20 Mbps and upload speeds of 3–10 Mbps. (The theoretical maximum for LTE is 300 Mbps download and 75 Mbps upload.)
- The range of 4G LTE depends on the tower and obstructions in the way. The optimal cell size is about 3.1 miles (5 km) in rural areas, and you can get reasonable performance for about 19 miles (30 km).

4G LTE was (and still is) pretty good for Internet access, provided the connection is relatively strong. Even though 5G has been in market for several years now and has good adoption, 4G is still being used and probably will be for several years to come. Many Internet of Things (IoT) systems use 4G because the throughput rates are good enough, latency is low, and the hardware needed is cheaper than 5G.

For anyone worried about 4G networks being shut down, think about this: 4G became available in 2008 and 5G earnestly started rolling out in 2019, yet cellular providers didn't start shutting down their 3G networks until 2022. According to some experts, 4G will probably be around until at least 2035.

Fifth Generation (5G)

The first *fifth generation (5G)* modem was announced in 2016, but it took until late 2018 for cellular providers to start test piloting 5G. At that point they had a substantial monetary investment in 4G and weren't yet sure of how to best implement the newer standard. 5G is fairly widespread but not available everywhere yet. Because we've already covered 5G in Volume 1 Chapter 9, we'll just review the highlights here:

- Some users will be able to get sustained wireless speeds in excess of 1 Gbps. The theoretical maximum peak download capacity is 20 Gbps.
- There are three 5G classifications:
  - Enhanced Mobile Broadband (eMBB), which is for cell phone and mobile communication
  - Ultra-Reliable Low-Latency Communications (URLLC) for autonomous vehicles and industrial applications
  - Massive Machine Type Communications (mMTC) for sensors, which is designed to support IoT devices

- The early focus was on developing eMBB infrastructure and mobile devices. Two versions of eMBB are available: fixed wireless broadband in densely populated areas and LTE for everywhere else.
  - 5G LTE is like 4G LTE but about 7–10 times faster. Browsing speeds for 5G are about 490 Mbps on average, and most users can get 100 Mbps downloads.
  - 5G uses existing LTE frequencies in the 600 MHz to 6 GHz range.
  - eMBB fixed wireless is revolutionary, offering gigabit speeds using millimeter wave (mmWave) technology in the 24 GHz to 86 GHz range.
  - mmWave is limited to about half a mile and is easily blocked by obstructions.

Having the ability to get gigabit performance over a cellular connection is exciting. Right now, it's hard to imagine how anyone could want or need more, but technology will continue to evolve. Ten or 20 years from now we'll look back and laugh at standards that only gave us a paltry gigabit per second and wonder who was able to survive on speeds that slow.

For now, though, gigabit through 5G is as good as it gets. Now that you're familiar with the standards, let's take a look at how to set up cellular data connections on mobile devices.

Using Cellular Data Connections

Nearly every user of a mobile device knows how to get on the Internet. iOS users have the built-in Safari browser, and Android users have Google Chrome available. Getting online when you have a cellular connection is easy to do if the device has a data plan. Another great feature of mobile devices is that they can share their cellular data connections with other devices. It's basically the exact opposite of joining a mobile phone to a Wi-Fi network, which we will talk about later in this chapter. Next, we will provide details on using and configuring hotspots and tethering, using airplane mode, and how data network updates are handled on mobile devices. We'll also cover a few key acronyms you need to know.

Using Mobile Hotspots and Tethering

We introduced mobile hotspots in Volume 1 Chapter 4, "Laptop and Mobile Device Hardware," and as an exam objective it's a topic that appears in a few places. Because it's an important concept for mobile device connectivity, and it's included in Exam Objective 1.3, we'll talk about how to use them here.

Real World Scenario

Where Is That Setting?

Throughout this chapter, we will use iOS 18.1.1 and Android 13 One UI Core 5.1 as example OSs for screenshots. Recognize that new versions of iOS and Android come out nearly every year, and the locations of features and settings for new OSs may change slightly. The point of showing screenshots in examples isn't for you to memorize exactly where

everything is, but rather to give you familiarity enough to accomplish the tasks you'll need to do.

Also note that in 2019, Apple transitioned its iPads from using iOS to using iPadOS. Throughout this chapter we will use the term iOS to refer to both generically, for brevity, unless we need to call out a specific option on a tablet versus a smartphone. Much like with version changes, options between iOS and iPadOS might differ slightly, but the option you're looking for will be there . . . somewhere.

In the real world, if the setting isn't exactly where you expect it to be, or if you're not very familiar with the OS you're working on, you can always Google the answer. It's an option that works for almost every hardware or software problem you run across.

Unfortunately, you can't use Google while sitting for the A+ exam. The good news is, CompTIA isn't going to test you on the differences between iOS 17.7.2 and iOS 18.1.1. What they may do is give you a screenshot of a settings page and ask you to configure something such as setting up corporate email. By understanding the concept of what you need to do—for example, toggling on email, setting the right protocol, and entering the email server name or address—you should be able to get the answer right. The example screen they give you might not be identical to what we showed here, but it should be close enough that you'll be able to accomplish the task with relative ease.

Recall that a *mobile hotspot* lets a device share its cellular Internet connection with Wi-Fi-capable devices. A Wi-Fi-enabled laptop, for example, would look for the mobile phone's Wi-Fi network, join it, and then have Internet access. To enable an iPhone to be a mobile hotspot, go to Settings ➢ Personal Hotspot. Figure 10.2 shows the personal hotspot screen. Simply slide the toggle to the On position to enable it. A password to join the network is provided (and can be changed) as well as instructions on how to join.

If the wireless plan does not allow that device to be a hotspot, then Settings ➢ Personal Hotspot will not exist. Instead, go to Settings ➢ Cellular. Tap the Set Up Personal Hotspot link to begin the process of setting it up.

Also recall that there are three potential challenges with using a smartphone as a mobile hotspot: speed, cost, and security. Cellular connections are usually slower than Wi-Fi, so having multiple devices trying to get on the Internet via one cellular link can be slow. From a cost standpoint, you could go over your data plan quite easily by using a hotspot, or the provider might charge you extra just to use it. Finally, there is security. iOS 7 and newer use WPA2, and the iPhone 11 and newer support WPA3, so security is less of an issue, but wireless networks are inherently unsecure due to the fact that the signals are transmitted through the air.

On the Android OS, a mobile hotspot is enabled under Settings ➢ Connections ➢ Mobile Hotspot And Tethering (see Figure 10.3). On Android, when you turn mobile hotspot on, it automatically turns Wi-Fi off.

FIGURE 10.2 Personal hotspot screen in iOS

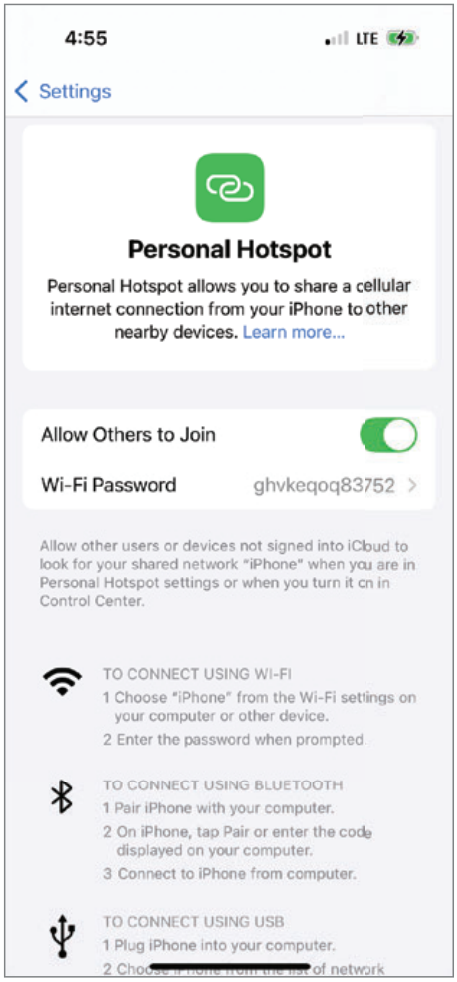
4:55
LTE
Settings
Personal Hotspot
Personal Hotspot allows you to share a cellular internet connection from your iPhone to other nearby devices. Learn more...
Allow Others to Join
Wi-Fi Password
ghvkeqoq83752
Allow other users or devices not signed into iCloud to look for your shared network "iPhone" when you are in Personal Hotspot settings or when you turn it on in Control Center.
TO CONNECT USING WI-FI
1 Choose "iPhone" from the Wi-Fi settings on your computer or other device.
2 Enter the password when prompted
TO CONNECT USING BLUETOOTH
1 Pair iPhone with your computer.
2 On iPhone, tap Pair or enter the code displayed on your computer.
3 Connect to iPhone from computer.
TO CONNECT USING USB
1 Plug iPhone into your computer.

Tapping Mobile Hotspot will show the network name (SSID) and hotspot password, and tapping Advanced shows options and additional settings (Figure 10.4). Figure 10.4 also shows the security choice—Android options are WPA2-Personal, WPA3-Personal, and Open.

In Android, a mobile hotspot uses WPA2 encryption by default, but encryption can be set to Open so the device is an open portal. We highly recommend that you always leave the encryption on! You can also change the network name and password associated with the hotspot.

Some mobile providers limit the number of devices that can join the mobile hotspot. For example, Verizon limits it to 10 devices for 4G LTE phones, and the Android phone used in the example also allows a maximum of 10 devices.

FIGURE 10.3 Mobile hotspot in Android

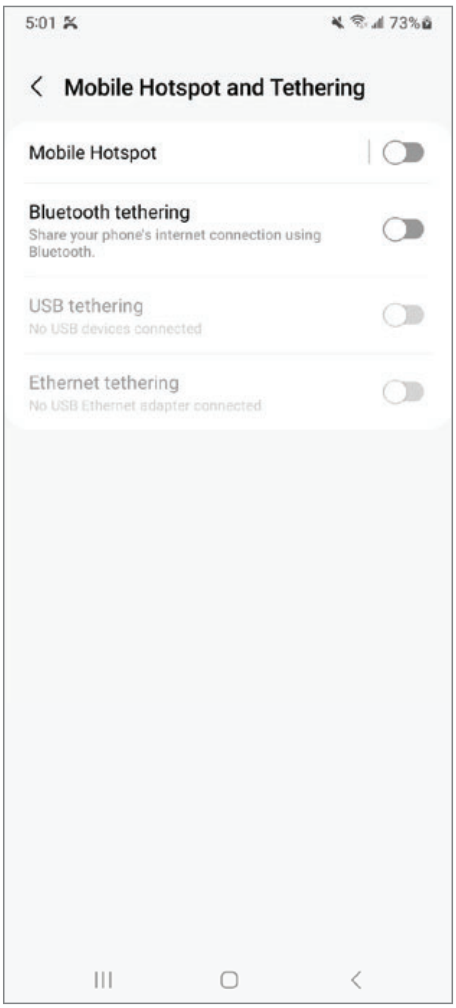

Finally, mobile providers sell small devices that are specifically used as mobile hotspots. Figure 10.5 shows an example of a Verizon Wireless MiFi hotspot. These types of devices will either use your existing mobile contract or will need to have an activation of their own.

Tethering is when you have connected a device to a mobile hotspot. The term used to be reserved only for when you connected via USB cable, as opposed to connecting via wireless. Some devices will not function as a mobile hotspot but will allow you to tether a laptop (or other device) to it so that the mobile device can share the cellular Internet connection.

Enabling and Disabling Wireless/Cellular Connections

Each type of wireless connection can be individually enabled or disabled under Settings. For example, Figure 10.6 shows the connections in iOS Settings and Android Settings ➢ Connections. Turning off one connection at a time serves its purpose—for example, disabling

Bluetooth to save battery life—but there's also an option to disable multiple wireless connections at once. It's called airplane mode.

FIGURE 10.4 Android hotspot network name, password, and options

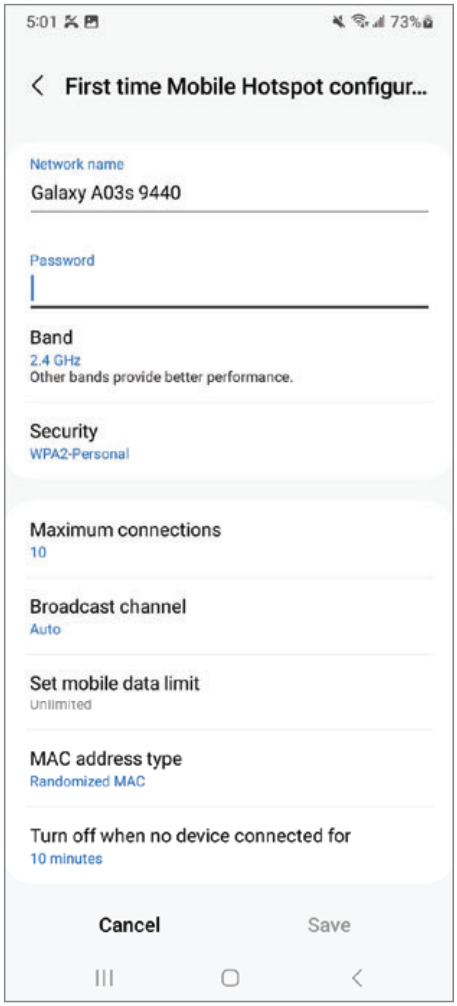

The *airplane mode* feature was named so because, for many years, no network signals were allowed on airplanes. Today, some airlines allow for in-flight Wi-Fi (for a nominal fee, of course), and Bluetooth is permitted, but the name of the feature still sticks. It's not restricted to airplane use, though. If you're in a public area and suspect that someone is trying to hack your phone through the Wi-Fi or Bluetooth connection, airplane mode will quickly shut down all your external connections. Android and iOS both make it easy to enable airplane mode and give you a few ways to get to it.

FIGURE 10.5 Verizon Wireless MiFi hotspot

FIGURE 10.6 Connections in (a) iOS and (b) Android

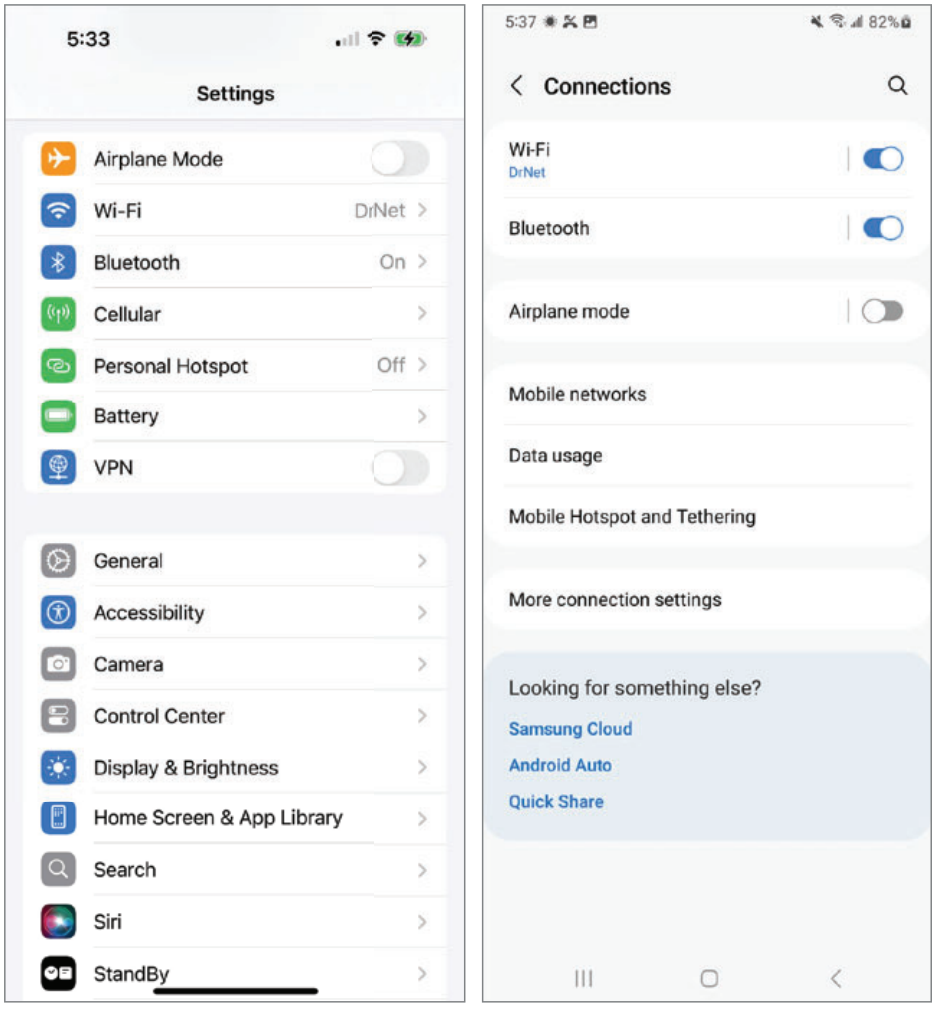

There are two quick ways to enable airplane mode in each OS. If you look at Figure 10.6, you can see that each OS has an airplane icon on its connections page. The second is to swipe down from the top of the screen. In iOS this opens the Control Center, and in Android it opens the notification area, shown in Figure 10.7. (To be exact, the Android screen shot in Figure 10.7 is when you swipe down twice to see the full notification area. Swiping down once gives a limited version, and the airplane mode icon is there too.) Tap the airplane icon to enable airplane mode.

FIGURE 10.7 (a) Control center in iOS and (b) Android notification area

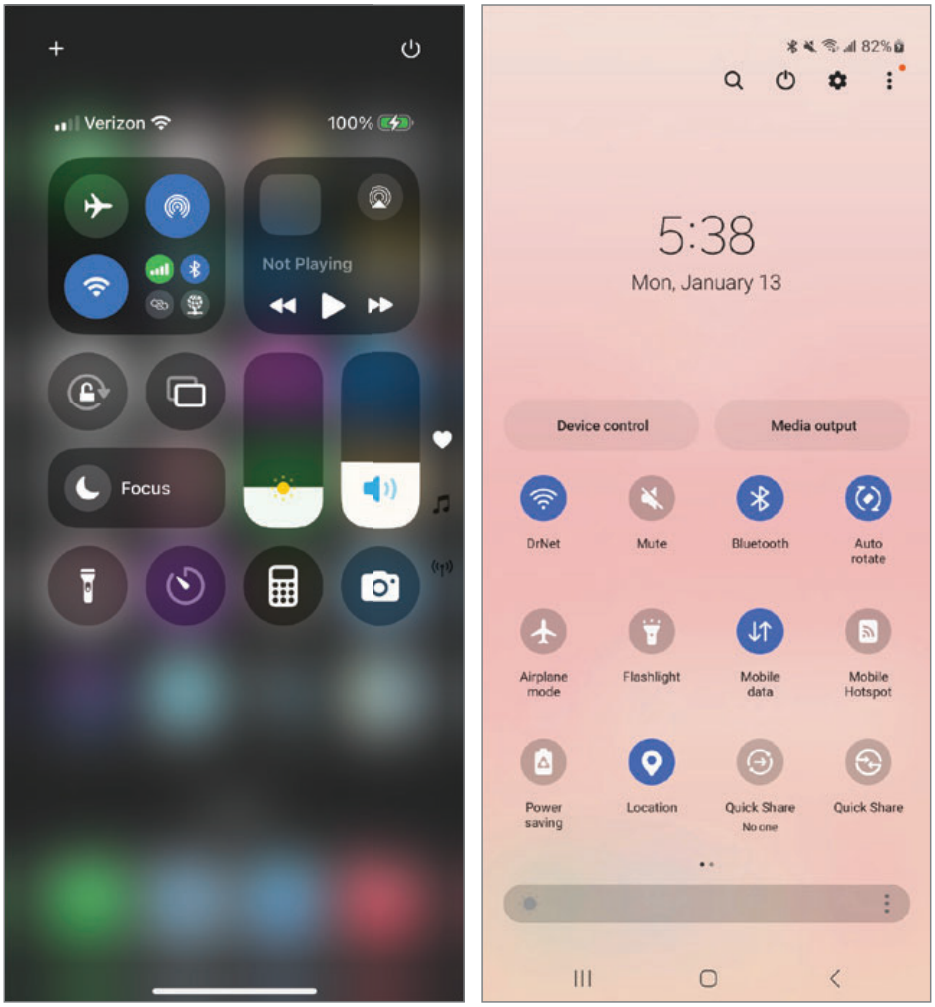

In newer versions of iOS, enabling airplane might not turn off all connections by default. Airplane mode remembers its settings from the last time it was enabled and enacts those. So, if you want to have airplane mode disable cellular and Wi-Fi, but leave Bluetooth enabled, first, enter airplane mode and configure it that way. Then, toggle airplane mode off. When you turn it on again, it will be configured as it was last time you were in it.

Understanding Cellular Updates

When most people think of cellular updates, they probably think of an update to the operating system. Perhaps a new Android version is available or iOS is alerting them to download its latest incarnation. Those updates are normal and completing them takes the active participation of the user. Other updates can occur too, and many of these are transparent to the user.

Before we talk about what those updates are, though, you must first understand that mobile phones don't just have one operating system. This might come as a surprise, but most mobile phones have three operating systems. Duties are split up among the operating systems, simply because there are so many specialized tasks for the phone to perform.

The first OS is pretty obvious—it's the one you interact with. The other two are specialized OSs that handle specific functions for the device. These two OSs are very small, typically only a few hundred kilobytes in size, and they are referred to as *real-time operating systems (RTOSs)*. They are designed to be lightweight and fast, and "real-time" refers to their ability to minimize lag in data transfers.

First, there is a *baseband OS* that manages all wireless communication, which is handled by a separate processor. Some people call the wireless communications chips in a mobile phone the radio or the modem. Consequently, you might hear about *radio firmware*, a *radio firmware update*, or a *modem update*. The last two terms are interchangeable with *baseband update*, which simply means an update of the baseband OS.

Real World Scenario

Why Phones Need a Baseband OS

One question many students ask is, "Why does a mobile phone need three OSs?" It would seem that phones perform fewer functions than PCs perform, and PCs run on one OS.

The answer is that it increases flexibility and reduces software development issues. Think about how many different types of mobile devices are on the market, and then think about how many versions of each are available. We don't know the exact answer, but it's a large number. Without the baseband OS, the primary OS would need to know how to talk to every single type of radio hardware on the market. And when new hardware comes out, which happens frequently, the primary OS would need to know how to talk to it as well. This isn't feasible. Therefore, phones were designed with one additional layer, which is the baseband OS. The baseband OS knows how to talk to the radio hardware, but it also knows how to talk to the primary OS. When a new hardware chipset is deployed, the developers just need to make the new baseband OS talk to iOS or Android on one side and the new hardware on the other side.

Desktop OSs face similar issues, but Windows, macOS, Linux, and others deal with this by using software drivers.

Second, a *subscriber identity module (SIM) OS* manages all data transfers between the phone and the *subscriber identity module (SIM) chip*, which is a small memory chip that stores user account information, phone identification, and security data, and it is generally tied to a specific carrier. An Android example is shown in Figure 10.8.

FIGURE 10.8 Android SIM chip

The SIM chip in Figure 10.8 is removable. Many manufacturers are now integrating SIM functions onto the smartphone's motherboard and calling it an *eSIM*. There are advantages and disadvantages to this system. Proponents say it's better because the SIM chip can't get stolen, whereas opponents say it makes it harder for users to change networks. We're not here to take a side in this debate but know that the eSIM chip can be updated much like flashing the BIOS on a PC.

SIM and eSIM are part of Exam Objective 1.3. Understand what they do and how they differ from each other.

These RTOSs are normally updated when a user updates an operating system, but occasionally the carrier will update them when the phone is not otherwise busy. Apple currently provides no way to update either RTOS manually on iOS devices. (Users can find information on how to jailbreak the phone online, but that voids all warranties and is not recommended.) There's more information available on how to update an RTOS on Android phones because Android is open source. Users or companies can provide newer versions of the baseband RTOS, and others can download and install them. Some will say that updating your baseband firmware can result in better reception, faster data throughput, and reduced battery usage. There is much Internet debate about the rewards versus the risks, though.

A discussion of performing a manual update of an RTOS requires three warnings:

- Be sure that the upgrade comes from a reputable source. Look for reviews or comments to be sure that it delivers what you expect it to provide.

- Understand that performing this update will most likely void any warranty that you have.
- Always be sure to allow the updates to complete. Interrupting a firmware update in the middle is a nearly surefire way to *brick* your phone—that is, to make it inoperable.

Two other updates of which you should be aware are *product release instruction (PRI) updates* and *preferred roaming list (PRL) updates*. The PRI contains settings for configuration items on the device that are specific to the network that it's on. The PRL is the reference guide the phone uses to connect to the proper cell phone tower when roaming. Both PRI updates and PRL updates also normally happen when the primary OS on the phone is updated. Some carriers make these two easier to update manually than the RTOS on Android phones, though. For example, Verizon users can dial **`*228`** for a manual PRL update. As always, check with the carrier before attempting to perform these updates and to determine the exact procedure.

Cellular Acronyms to Know

The last section introduced a few new acronyms to know, such as PRI and PRL. There are a few others you might see in relation to mobile phones that you should know too:

International Mobile Equipment Identity The *International Mobile Equipment Identity (IMEI)* is a 15-digit serial number that is unique to each phone. If a phone is reported stolen, the IMEI will be declared invalid and the phone disabled. The IMEI can be displayed on most phones by dialing **`*#06#`**. AT&T and T-Mobile were the first networks to use IMEI.

Mobile Equipment Identifier The *mobile equipment identifier (MEID)* is an alternate form of a serial number. It's identical to the first 14 numbers of the IMEI. Sprint (now T-Mobile) and Verizon were the first to use MEIDs.

International Mobile Subscriber Identity The *international mobile subscriber identity (IMSI)* is a unique 15-digit identifier that describes a specific mobile user and their network. It's composed of the following three elements:

Mobile Country Code The *mobile country code (MCC)* is a three-digit code, such as 310 for the United States and 234 for the United Kingdom.

Mobile Network Code The *mobile network code (MNC)* is a two- or three-digit code that identifies the carrier. Many carriers have multiple codes. For example, in the United States, 006 is a Verizon code, and 170 and 410 (among others) are AT&T codes.

Mobile Station Identifier Number The mobile station identifier number (MSIN) is a sequential serial number.

Integrated Circuit Card Identifier The *integrated circuit card identifier (ICCID)* is a 19- or 20-digit identifier for each SIM chip globally. It's like a serial number for the SIM card.

Secure Element Identifier The *secure element identifier (SEID)* is a very long hexadecimal code that uniquely identifies the phone. It is used in security applications, NFC, and features like Apple Pay.

Within iOS, you can find many of these numbers by choosing Settings ➢ General ➢ About and scrolling to the bottom, as shown in Figure 10.9. To find the same information in Android, go to Settings ➢ About Phone (also Figure 10.9). You might need the IMEI or other information if you are troubleshooting connectivity issues with a cellular provider.

FIGURE 10.9 (a) About in iOS and (b) About Phone in Android

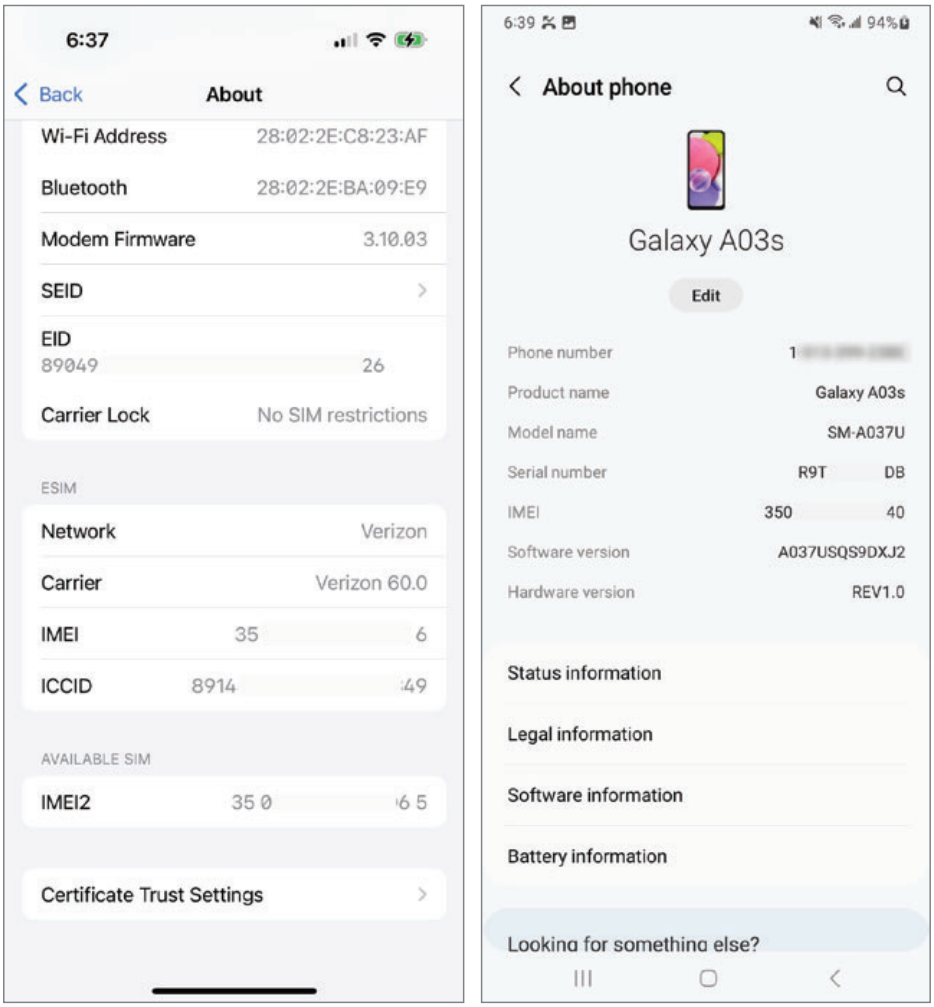

Establishing Wi-Fi Connectivity

Using a cellular network is great because you can connect from nearly anywhere. The downsides, though, are that the connection is slower than connectivity methods such as Wi-Fi, and

you have to pay for the data you use. When within range of a secured Wi-Fi network, take advantage of the device's ability to use that network instead. Not only will the connection be faster, it will be free.

Being able to configure Wi-Fi connectivity is a key component of Exam Objective 1.3. Understand how to enable Wi-Fi connectivity and access a network in both Android and iOS.

Before you can transfer data over a Wi-Fi network, you have to attach to the network in the same manner you would attach a laptop, for instance, to the same wireless network. You have to find the network by its service-set identifier (SSID), or you have to enter the SSID if it is not being broadcast. You must then satisfy any security requirements that might be in place, such as using WPA3 or having the right security keys. Exercise 10.1 steps you through the procedure on an iPhone.

EXERCISE 10.1

Connecting an iPhone to a Wi-Fi Network

1. Tap the Settings app on the Home screen.
2. Tap Wi-Fi on the Settings menu.
3. On the Wi-Fi menu, swipe the Wi-Fi switch to the right to turn it on, as shown in Figure 10.10.

FIGURE 10.10 Turning on Wi-Fi in iOS

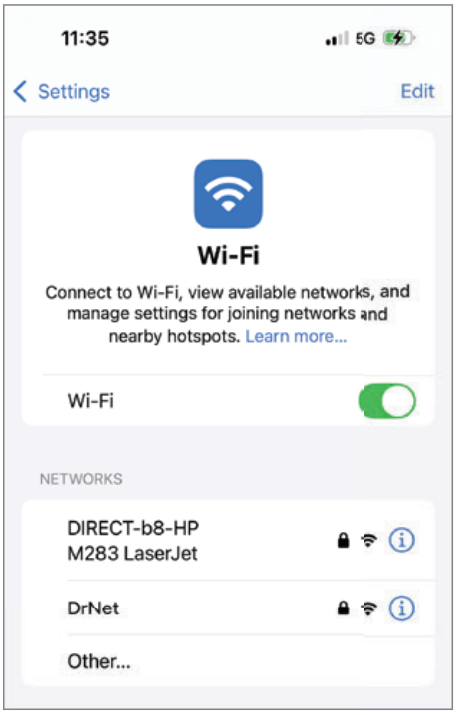

Tapping switches also works to toggle them to the opposite state.

4. In the Networks list, tap the name of the wireless network that you want to join.

   If the network you want to join is not listed, tap Other and you will be asked to enter the network name you want to join.

5. Enter the password or key for the wireless network, if you are asked for one, and then tap the Join button.

   If it connects, it will move the network name to under the Wi-Fi switch and put a check next to it, as shown in Figure 10.11. Also note the Ask To Join Networks option near the bottom. Tapping it gives you the options of Off, Notify, and Ask. Turning this off means that you will never be interrupted with offers to join new wireless networks that are in range of the device. Regardless of the setting of this switch, the device will still automatically reconnect to remembered Wi-Fi networks as they come into range. Otherwise, it will join if the other network has a stronger signal and then notify you (Notify) or ask if it should join the new network (Ask).

FIGURE 10.11 Wi-Fi settings with a network that's been joined

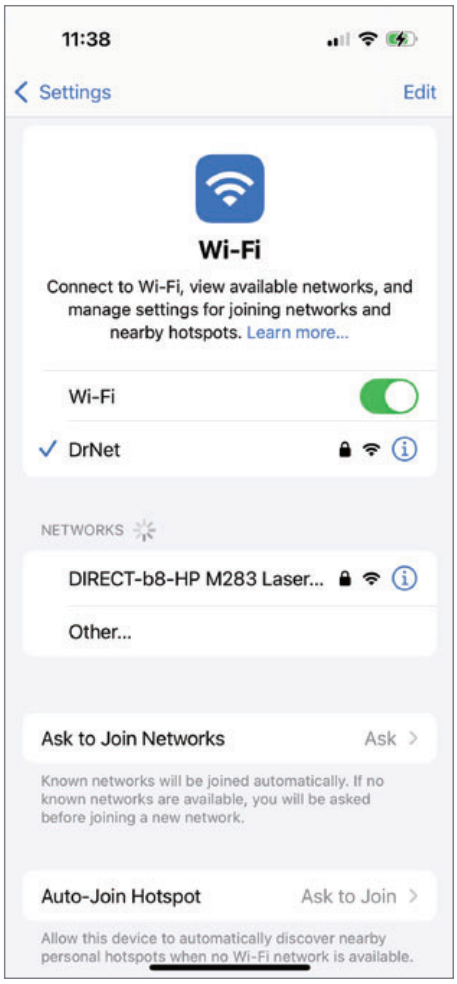
11:38
Settings
Edit
Wi-Fi
Connect to Wi-Fi, view available networks, and manage settings for joining networks and nearby hotspots. Learn more...
Wi-Fi
DrNet
NETWORKS
DIRECT-b8-HP M283 Laser...
Other...
Ask to Join Networks
Ask
Known networks will be joined automatically. If no known networks are available, you will be asked before joining a new network.
Auto-Join Hotspot
Ask to Join
Allow this device to automatically discover nearby personal hotspots when no Wi-Fi network is available.

EXERCISE 10.1 *(continued)*

6. Tap the Settings back button at the top left to return to the previous screen.

   Notice on this screen that the network you selected is now listed on the Settings page, as shown in Figure 10.12. If not, you are not connected to a Wi-Fi network.

FIGURE 10.12 Settings page with Wi-Fi network connected

7. Close Settings.

Throughout this chapter, the practice of tapping successive back buttons in the upper-left corner of the screen, instead of swiping up from the Home bar (or clicking the Home button on phones older than iPhone X), will be referred to as *backing out of the app*.

A similar series of tasks is required when attaching an Android phone to the same type of network. Exercise 10.2 details that procedure.

EXERCISE 10.2

Connecting an Android Phone to a Wi-Fi Network

1. Tap the Settings app to open it, and then tap Connections.
2. Slide the Wi-Fi switch to the On position, as shown in Figure 10.13.

FIGURE 10.13 Enabling Wi-Fi in Android

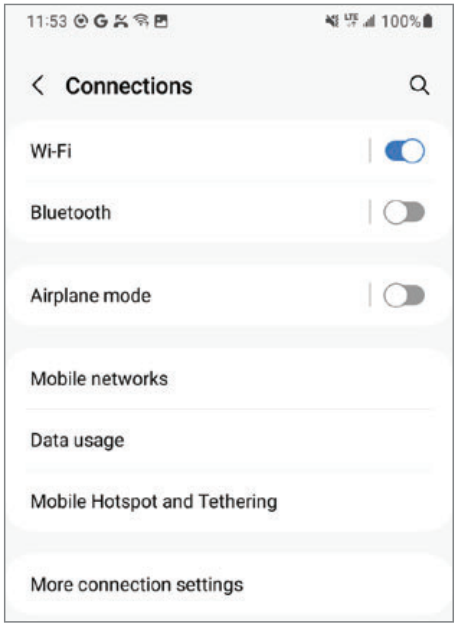

3. Tap Wi-Fi to see the list of available networks, as shown in Figure 10.14.

FIGURE 10.14 List of available networks

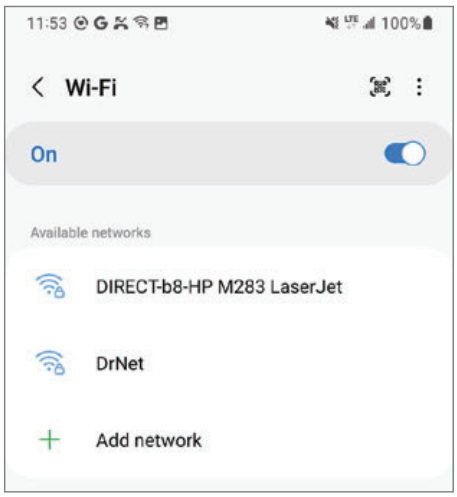

4. Tap the network you want to join.
5. Enter the password for the network and tap Connect. Also notice that on this screen there is a toggle switch to auto-reconnect.

EXERCISE 10.2 *(continued)*

The device will go back to the list of Wi-Fi networks, showing that it's connected (see Figure 10.15).

FIGURE 10.15 Connected to a network

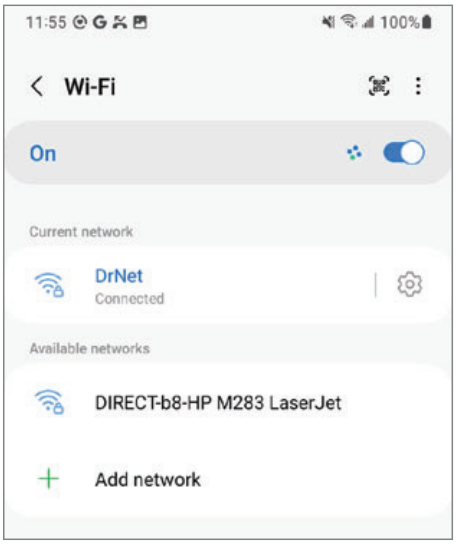

6. Close the Settings app.

When your phone is connected to a Wi-Fi network, you don't need to use a cellular connection for data transfers—apps will use the Wi-Fi connection for data. But if the connection gets dropped or you move out of Wi-Fi range, the device will use the cellular connection. This might be fine, but it also might not be what you want. If you want to ensure that the phone does not use cellular for data connections, you can disable that option. Exercise 10.3 walks you through the steps of how to do that on an iPhone. When the device is connected to a Wi-Fi network or when paired with a Bluetooth peer, data access will be possible; otherwise, no data-network access will occur.

EXERCISE 10.3

Disabling Cellular Use for Data Networking on an iPhone

1. Tap the Settings app on the home screen.
2. Select Cellular from the Settings menu.
3. Turn off the switch labeled Cellular Data, which is shown in the On position in Figure 10.16.

FIGURE 10.16 iOS Cellular settings

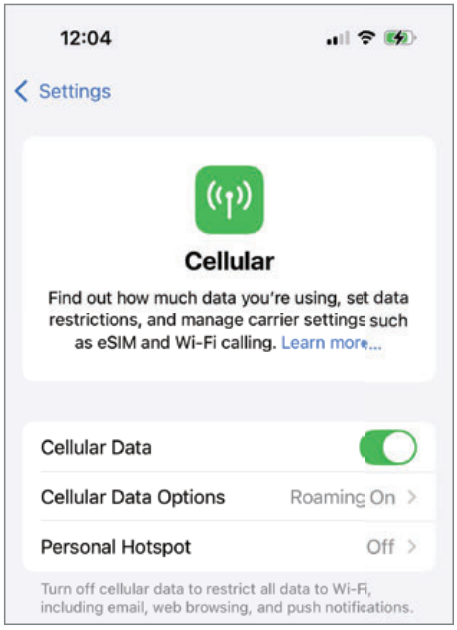
12:04
Settings
Cellular
Find out how much data you're using, set data restrictions, and manage carrier settings such as eSIM and Wi-Fi calling. Learn more...
Cellular Data
Cellular Data Options
Roaming On
Personal Hotspot
Off
Turn off cellular data to restrict all data to Wi-Fi, including email, web browsing, and push notifications.

4. If you would like to keep cellular data usage enabled but not allow roaming into other providers' data networks, you can tap Cellular Data Options, which will bring up the options shown in Figure 10.17.

FIGURE 10.17 Cellular data options

5. Here you have the options to disable voice and/or data roaming, as well as configure how voice and data is handled, such as whether to use 5G or LTE. Setting Data Mode to Low Data Mode can help reduce cellular data usage.

EXERCISE 10.3 *(continued)*

6. If you wish to disable cellular use just for a specific function or app, scroll down in the Cellular settings. You can turn off the use of cellular data for that function or app alone.

   For example, Figure 10.18 shows the Cellular Data Usage page for several apps, all of which are currently enabled. Tap the app name to disable cellular usage for that app. This page also shows you how much cellular data each app has used in the current period.

FIGURE 10.18 Cellular data for individual apps

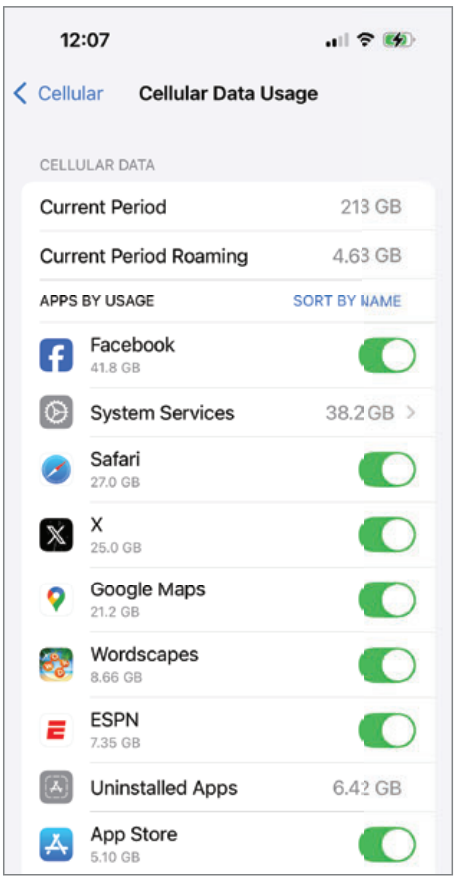

If the app you're looking for isn't on this page, you can also disable cellular data for an app by going to Settings ➢ Apps, finding the app in the list, and tapping on it. Figure 10.19 shows an example, and here you can see the Cellular Data option that you can toggle off or on.

FIGURE 10.19 App settings

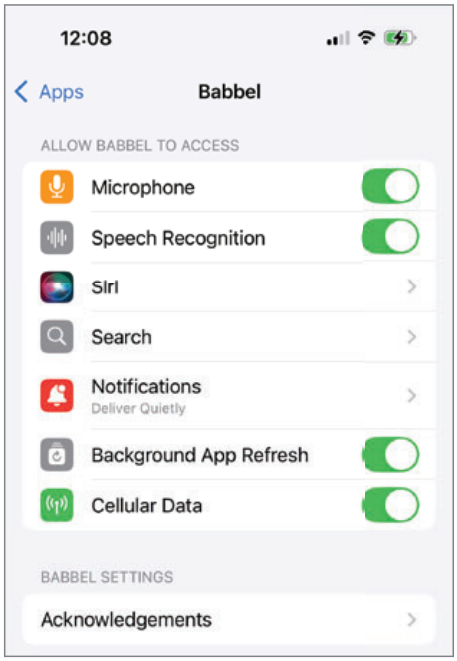

7. Back out of the Settings app or close it.

Exercise 10.4 shows you how to disable cellular data on an Android-based device.

EXERCISE 10.4

Disabling Cellular Use for Data Networking in Android OS

1. Open the Settings app, and then tap Connections.
2. Tap Data Usage.

   It will bring up a screen similar to the one shown in Figure 10.20.

EXERCISE 10.4 *(continued)*

FIGURE 10.20 Data usage in Android OS

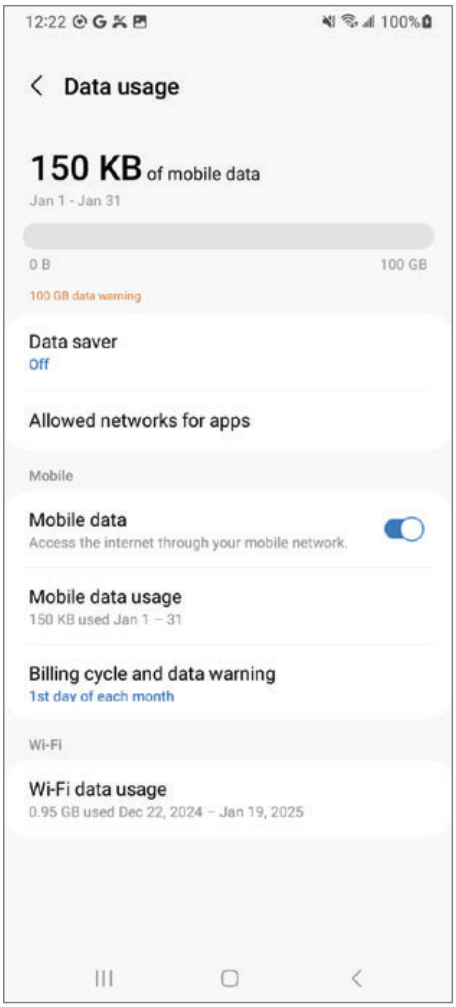

3. To disable cellular data usage, slide the Mobile Data switch to the Off position.

   Note that you can leave cellular data enabled but limit the amount that can be used during a specified time period. The dates can be customized to match the billing period used by the mobile carrier. However, be warned that the carrier may track usage slightly differently from the phone, so it might not be a foolproof way to avoid excess data charges.

   Another handy option on this screen is Data Saver. When turned on, it reduces the background data any app can use, essentially forcing all apps to use Wi-Fi only unless

otherwise specified in the app's settings. In other words, this helps reduce cellular data usage.

4. (Optional) To see how much data each app has used in the current period, scroll down and tap Wi-Fi data usage (see Figure 10.21).

FIGURE 10.21 Viewing Wi-Fi app data usage in Android OS

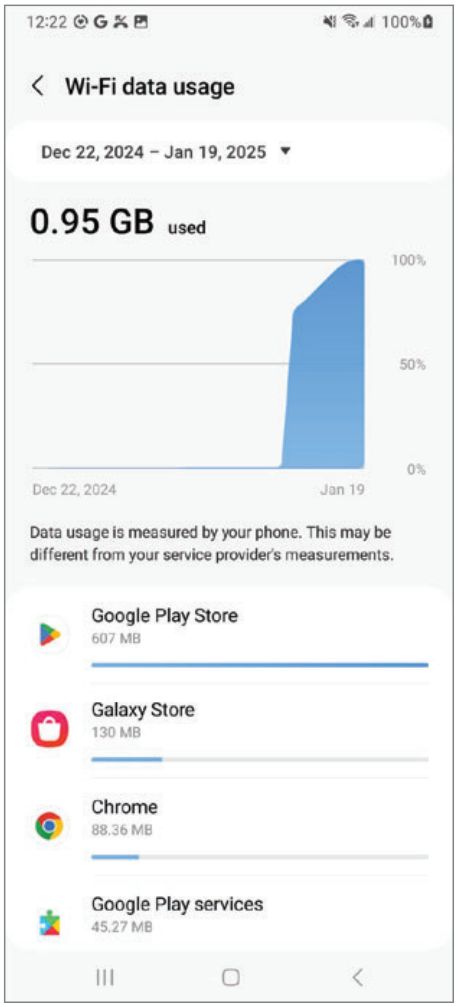

5. (Optional) If you tap on an app in the list, you can see its data usage configuration options (Figure 10.22). This is where you can configure it to still use cellular connections when Data Saver is on. You can also get to this screen by going to Settings ➢ Apps, tapping the app name, and then tapping Mobile Data.

EXERCISE 10.4 *(continued)*

FIGURE 10.22 Individual application data usage

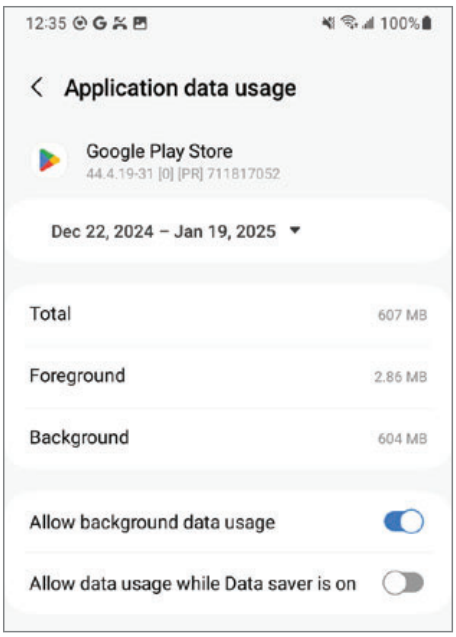

The final setting we will look at in relation to Wi-Fi networks is the *virtual private network (VPN)* configuration. Recall from Volume 1 Chapter 7, "Introduction to TCP/IP," that a VPN is a secured network connection made over an unsecure network, such as the Internet. For example, if you wanted to connect your phone to your corporate network over the Internet in order to read email, but you also wanted to secure the connection, you could use a VPN. To set up a VPN on an iPhone, perform the following steps:

1. Select Settings ➢ General ➢ VPN & Device Management ➢ VPN. (Note that if this device has previously connected to a VPN, the VPN can be enabled under the main screen of Settings. Refer back to Figure 10.12 to see the VPN toggle.)

   You will see a screen similar to the one shown in Figure 10.23. You can see that there is one VPN already configured on this device but that VPN is turned off.

2. To add a new VPN connection, tap Add VPN Configuration, which opens a screen like the one shown in Figure 10.24.
3. Choose the security protocol type (IKEv2, IPsec, or L2TP), then provide a server name and remote ID, an authentication method (username or certificate), and a password.
4. If going through a proxy server is required for the VPN, configure that at the bottom of the screen.

FIGURE 10.23 VPN settings

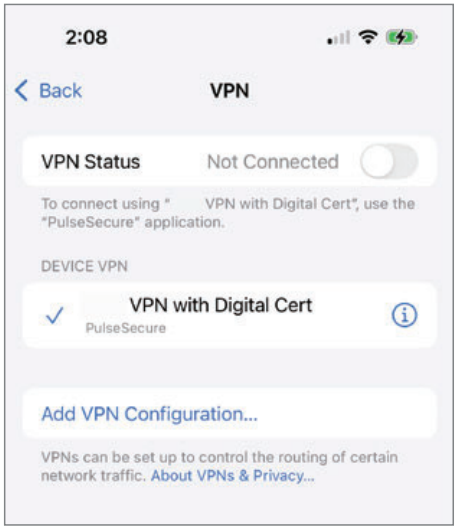
2:08
Back
VPN
VPN Status
Not Connected
To connect using " VPN with Digital Cert", use the "PulseSecure" application.
DEVICE VPN
VPN with Digital Cert
PulseSecure
Add VPN Configuration...
VPNs can be set up to control the routing of certain network traffic. About VPNs & Privacy...

FIGURE 10.24 Adding a VPN connection

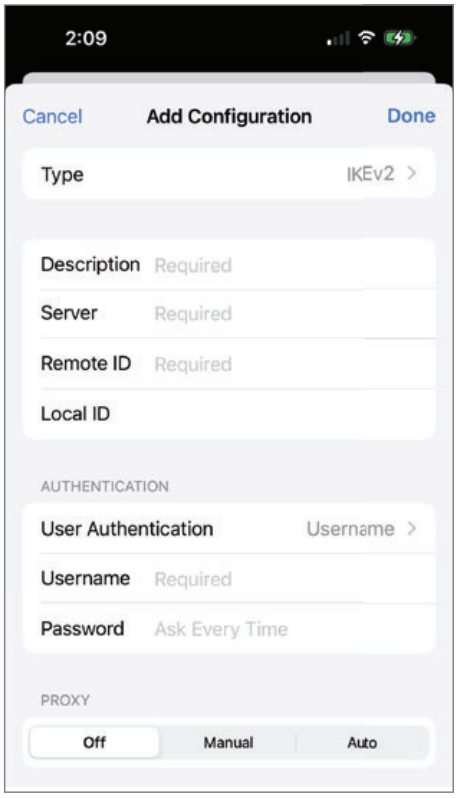

In order for the VPN to work, the server and client must be configured to use the same security protocols. iOS provides three VPN security protocol choices: Internet Key Exchange v2 (IKEv2), Internet Protocol Security (IPsec), and Layer 2 Tunneling Protocol (L2TP). IKEv2 and L2TP are both extensions of the IPsec protocol suite (which is in turn part of the TCP/IP suite you learned about in Volume 1 Chapter 7. IKEv2 is the most secure and the fastest, but it isn't supported by as many operating systems as L2TP.

Once you have enabled the VPN, a new VPN option will appear on your Settings page, as previously shown in Figure 10.23. This will allow you to easily enable, disable, or configure the VPN.

VPN connections can be made over cellular as well as over Wi-Fi.

Exercise 10.5 shows you the steps required to set up a PPTP or L2TP VPN connection in Android.

EXERCISE 10.5

Setting Up a VPN in Android

1. Open Settings ➢ Connections ➢ More Connection Settings (Figure 10.25). Tap VPN.

FIGURE 10.25 More connection settings

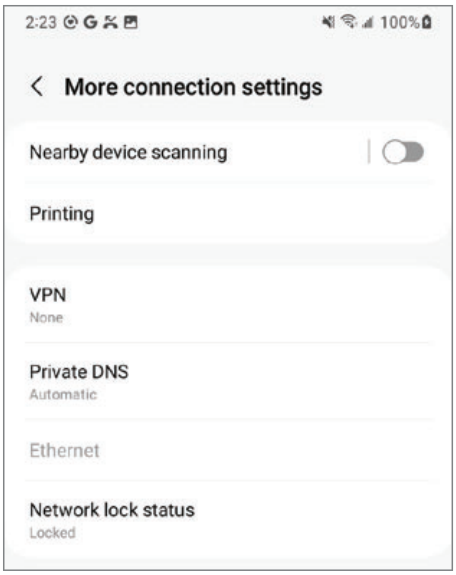

2. If there are no VPNs configured, it will say No VPNs. Tap the three dots icon in the upper-right corner and then tap Add VPN Profile. A screen similar to the one shown in Figure 10.26 will appear.

FIGURE 10.26 Edit VPN network

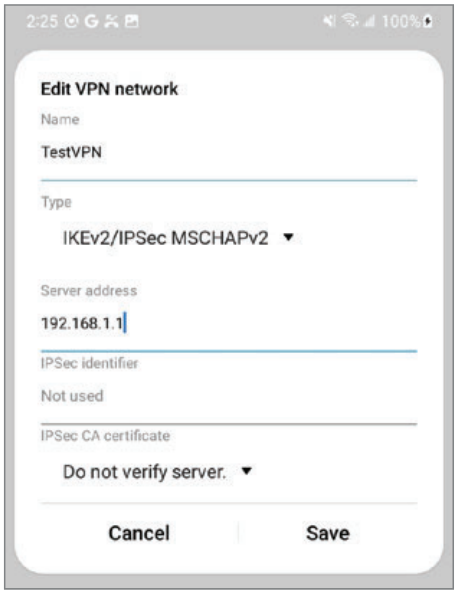

3. Enter the server's VPN configuration information, such as the name, security type (IKEv2/IPSec MSCHAPv2 or others—use the most secure method that the server requires), and address, as well as the security certificate information or username and password (depending on the security method chosen).
4. Tap Save. That will take you back to the Add VPN screen, with the new VPN listed (Figure 10.27).

FIGURE 10.27 Saved VPN

EXERCISE 10.5 *(continued)*

5. Tap the name of the VPN, which will open the Connect To VPN screen (Figure 10.28).

FIGURE 10.28 Connect to VPN

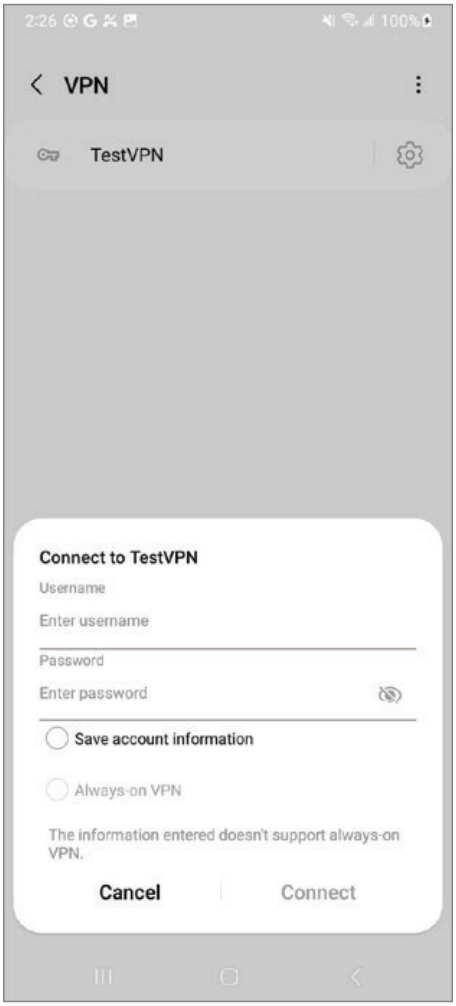

6. If required (it will depend on the security type chosen), enter the username and password, and tap Connect.

Mobile Phone Security

This is beyond what you need to know for the A+ exam, but you might find it interesting anyway. The security options on Android devices are likely different from iOS-based options. For example, the phone we use in Exercise 10.5 supports IKEv2/IPsec MSCHAP v2.

IKEv2 is the connection protocol, IPSec is the tunnelling protocol that handles encryption, and Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2) is the authentication protocol.

Android devices may also show options such as L2TP/IPsec PSK or L2TP/IPsec RSA. Pre-shared Key (PSK) is a security key that's generated by the server and shared with the client. It's susceptible to brute-force attacks and not recommended for highly secure environments. Rivest, Shamir, Adleman (RSA) is a somewhat more secure algorithm that uses asymmetric security keys. It's a better choice than PSK but needs to be administered from an RSA server.

Android also supports many apps that allow you to configure VPN connections, such as TunnelBear (`www.tunnelbear.com`) owned by McAfee and Hola Free VPN (`https://hola.org`).

A very common VPN implementation is called OpenVPN. Android and iOS do not natively support OpenVPN. If your network uses an OpenVPN server, know that in Android OS you have to install a third-party app (such as OpenVPN Connect) from the Google Play Store to create the VPN connection. iOS devices can download OpenVPN Connect from the App Store.

Establishing Bluetooth Connectivity

The IEEE 802.15 standard specifies wireless personal area networks (WPANs) that use Bluetooth for data-link transport. The concept is that certain paired devices will be capable of exchanging or synchronizing data over a Bluetooth connection, such as between a mobile device and a desktop or laptop computer.

In other cases, the Bluetooth *pairing*, or connecting two Bluetooth devices, can be used simply to control one device with another, allowing information to flow bidirectionally, even if that transfer does not result in its permanent storage on the destination. Examples of this latter functionality include a Bluetooth headset for a smartphone, a Bluetooth-attached keyboard and mouse, and pairing a smartphone or MP3 player with a vehicle's sound system.

In general, connecting a mobile device to another device requires that both devices have Bluetooth enabled. Pairing subsequently requires that at least one of the devices be discoverable and the other perform a search for Bluetooth devices. Once the device performing the search finds the other device, a sometime-configurable pairing code must often be entered on the device that performed the search. The code must match the one configured on the device that was found in order for the pairing to occur. In some cases, this pairing will work in one direction only. Usually, it is the mobile device that should search for the other device. If both devices have the same basic capability and will be able to exchange data readily, then it's not as important which device performs the search. Regardless, the pairing code must be known for entry into the device that requests it.

The truth about pairing mobile devices with conventional computers is that the results are hit or miss. There's never any guarantee that a given pairing will be successful to the point

of data transfer capability. Both devices must agree on the same Bluetooth specification. This turns out to be the easy part because devices negotiate during the connection. The part that is out of your control is what software services the manufacturer decided to include in their devices. If one device is not capable of file transfers over Bluetooth, then the pairing may go off without a hitch, but the communication process will stop there.

Exam Objective 1.3 states that you should be able to configure basic Bluetooth mobile device connectivity. This means you should know the steps required to pair a device. Per the exam objectives, those are as follows:

- Enable Bluetooth
- Enable pairing
- Find a device for pairing
- Enter the appropriate *personal identification number (PIN) code*
- Test connectivity

It would be wise to commit those steps to memory, as well as be able to execute a simple pairing if asked.

It sometimes takes a few tries to get the pairing or file transfer to work, so always be willing to try a few times. In the worst-case scenario, if it's still not working, look for documentation online to help. Exercise 10.6 shows the steps to connect an Android phone to a Windows 11 laptop over Bluetooth and then to transfer a file back and forth between the two. This exercise is split into three sections so that you can concentrate on individual stages of the pairing and file sharing processes.

EXERCISE 10.6

Pairing an Android Phone with a Windows Laptop

Enabling Bluetooth and Pairing

1. On the laptop, enable Bluetooth in Settings ➢ Bluetooth & Devices, as shown in Figure 10.29. (To get to Settings, tap the Start button and click the gear icon.)

FIGURE 10.29 Enabling Bluetooth in Windows 11

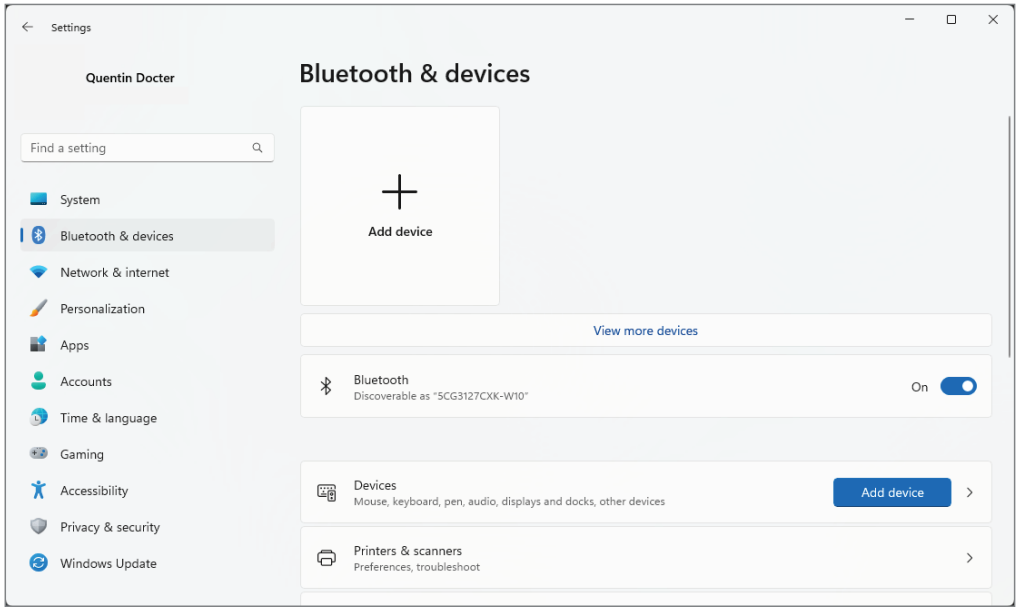

2. On the Android device, tap Settings ➤ Connections ➤ Bluetooth, and turn on Bluetooth.

   Under Available Devices (Figure 10.30), you should see the name of the laptop, provided you are within range. If no devices appear, tap Scan near the top to re-scan for devices.

FIGURE 10.30 Available Bluetooth devices

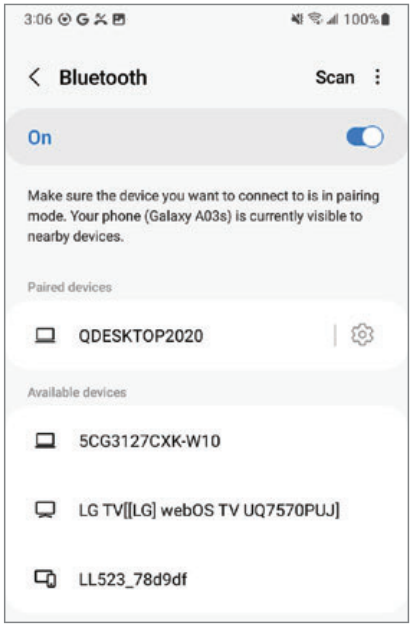

EXERCISE 10.6 *(continued)*

3. Tap the name of the laptop to connect to.

   It will begin pairing. A passkey will appear, as shown in Figure 10.31. Be sure the passkey matches the one shown on the laptop, and tap Pair. Note that you might also need to confirm the passkey on the laptop.

   Once the devices are paired, a gear icon will appear in Android, as shown in Figure 10.32.

FIGURE 10.31 Bluetooth pairing request

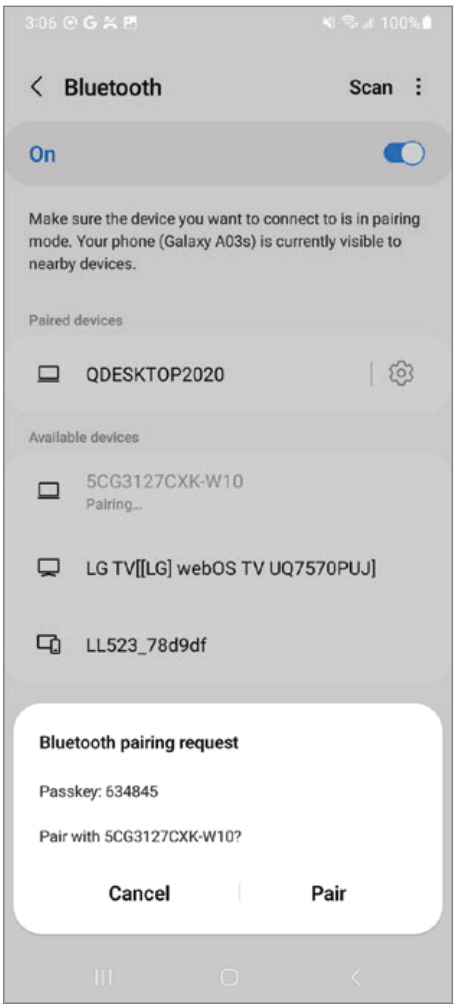

FIGURE 10.32 Paired Bluetooth device with gear

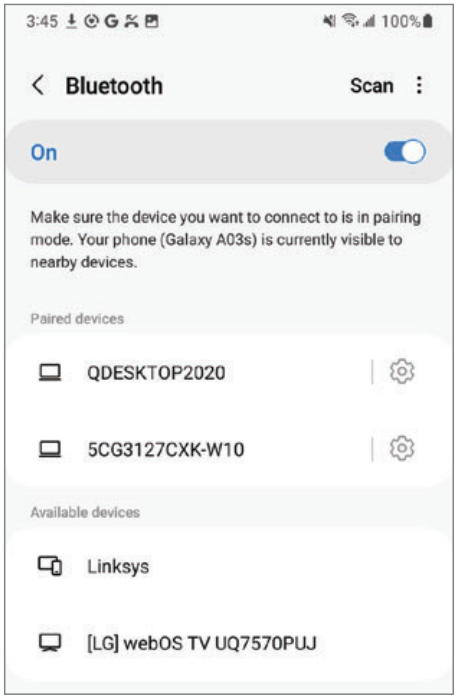

4. (Optional) To unpair the devices, tap on the gear and tap Unpair. You can also rename the paired device in this screen.

Transferring Files—Sending to Android

1. In Windows, open Settings ➢ Bluetooth & Devices.
2. Scroll down until you see the Send Or Receive Files Via Bluetooth option (Figure 10.33).

EXERCISE 10.6 *(continued)*

FIGURE 10.33 Bluetooth and Devices

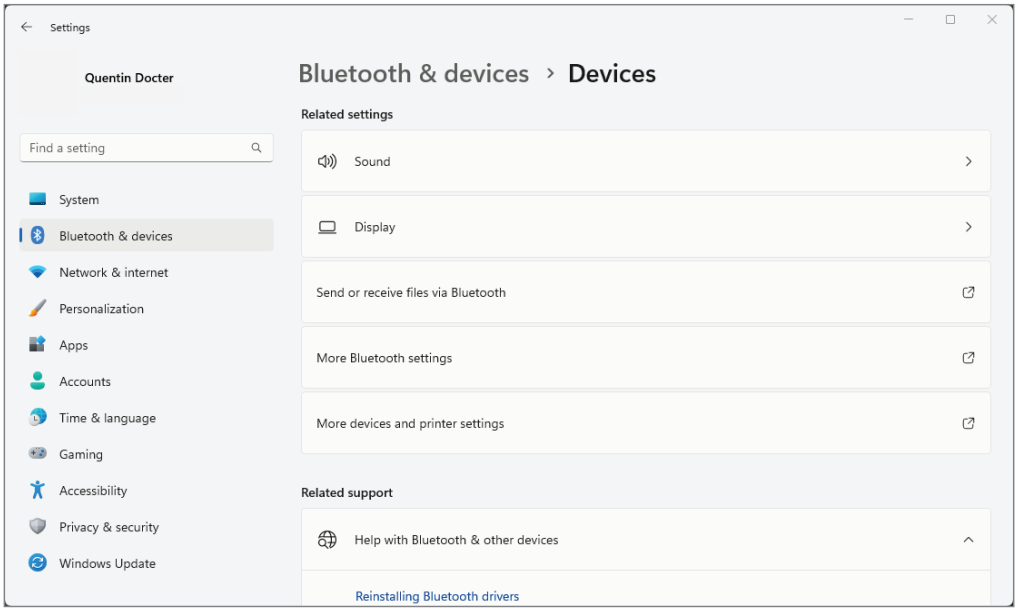

3. Click the Send Or Receive Files Via Bluetooth link.

   A screen will open (Figure 10.34), asking if you want to send or receive files.

FIGURE 10.34 Send or receive files via Bluetooth.

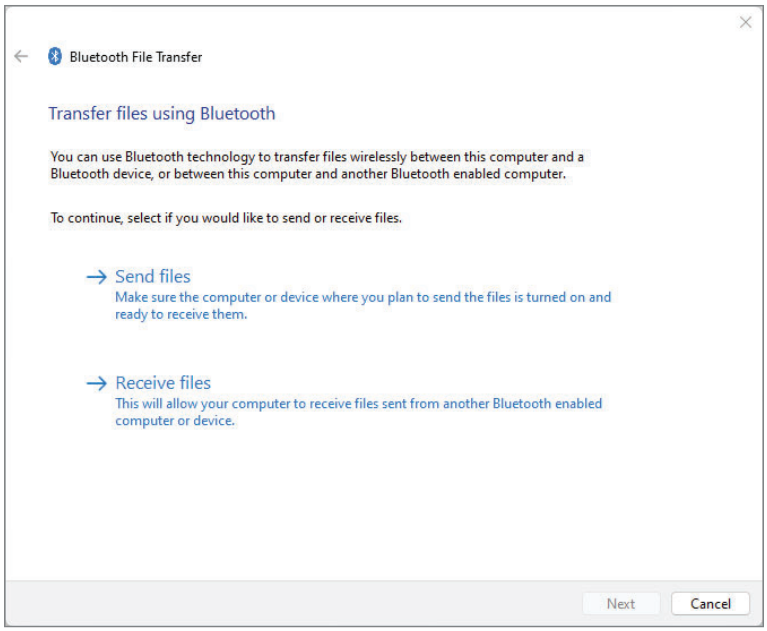
Bluetooth File Transfer
Transfer files using Bluetooth
You can use Bluetooth technology to transfer files wirelessly between this computer and a Bluetooth device, or between this computer and another Bluetooth enabled computer.
To continue, select if you would like to send or receive files.
Send files
Make sure the computer or device where you plan to send the files is turned on and ready to receive them.
Receive files
This will allow your computer to receive files sent from another Bluetooth enabled computer or device.
Next
Cancel

4. Click Send Files.
5. Choose the paired device (Figure 10.35) to send files to and click Next.

FIGURE 10.35 Choosing the paired device

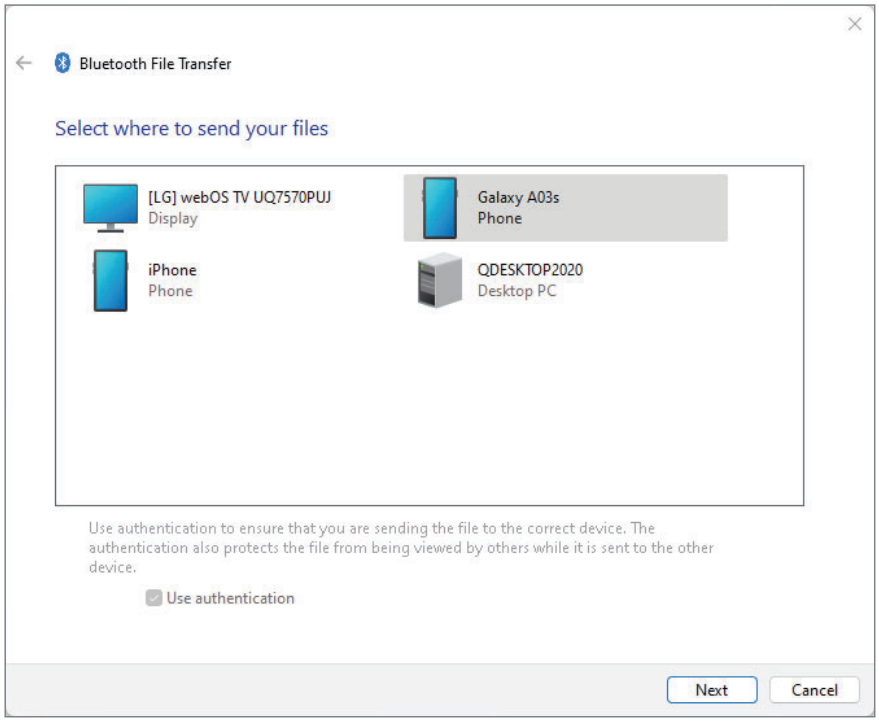
Bluetooth File Transfer
Select where to send your files
[LG] webOS TV UQ7570PUJ
Display
Galaxy A03s
Phone
iPhone
Phone
QDESKTOP2020
Desktop PC
Use authentication to ensure that you are sending the file to the correct device. The authentication also protects the file from being viewed by others while it is sent to the other device.
Use authentication
Next
Cancel

6. Browse for the files to send. You can send multiple files by holding down the Ctrl key while clicking files.
7. Click Next.
8. On the Android device, click the Accept button (Figure 10.36).

   On the laptop, there will be a transfer status bar. When the transfer is done, the laptop will say File Successfully Transferred.

EXERCISE 10.6 *(continued)*

FIGURE 10.36 Accepting Bluetooth file transfer

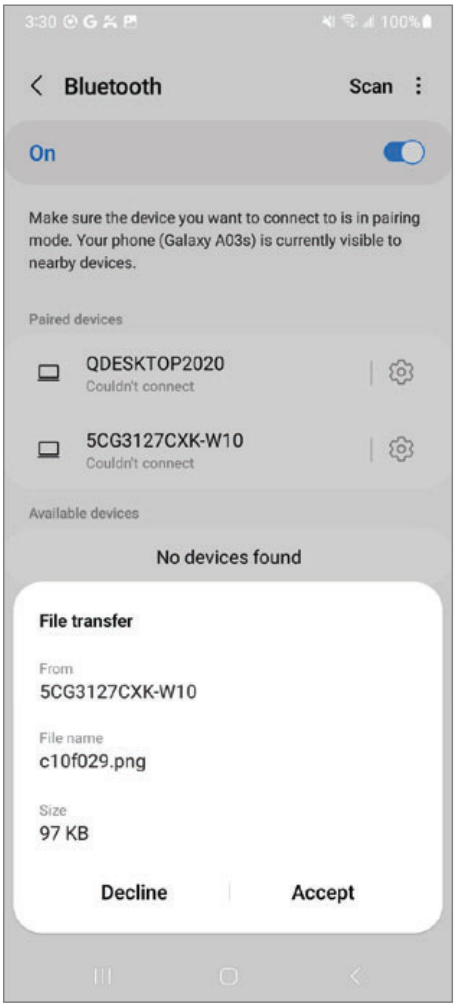

9. Click Finish.
10. In Android, open My Files ➢ Documents.

    The files will be in the `Download` folder by default (Figure 10.37).

FIGURE 10.37 Received file in Android

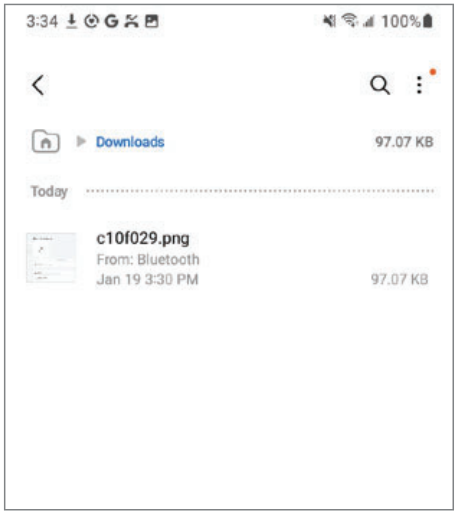

Transferring Files—Receiving from Android

1. In Windows, open Settings ➢ Bluetooth & Devices.
2. Scroll down until you see the option Send Or Receive Files Via Bluetooth.
3. Click or tap the Send Or Receive Files Via Bluetooth link.

   It will open a screen asking if you want to send or receive files.
4. Click Receive Files.

   The next window that appears will tell you the laptop is waiting for a connection.
5. On the Android device, find the file you want to share. Hold down on the filename to highlight it (tapping it just once will open up an app to view the file).
6. Tap the Share icon at the bottom of the screen (Figure 10.38, it's in the center—the three dots with the lines connecting them) and then Bluetooth.

EXERCISE 10.6 *(continued)*

FIGURE 10.38 Sharing a file

7. Tap the name of the device you want to share with.

   On the laptop, a status bar will appear as the file is being received.

8. The laptop will ask for a location to save the file. Choose a location, as shown in Figure 10.39.

FIGURE 10.39 Choose a location to save the file.

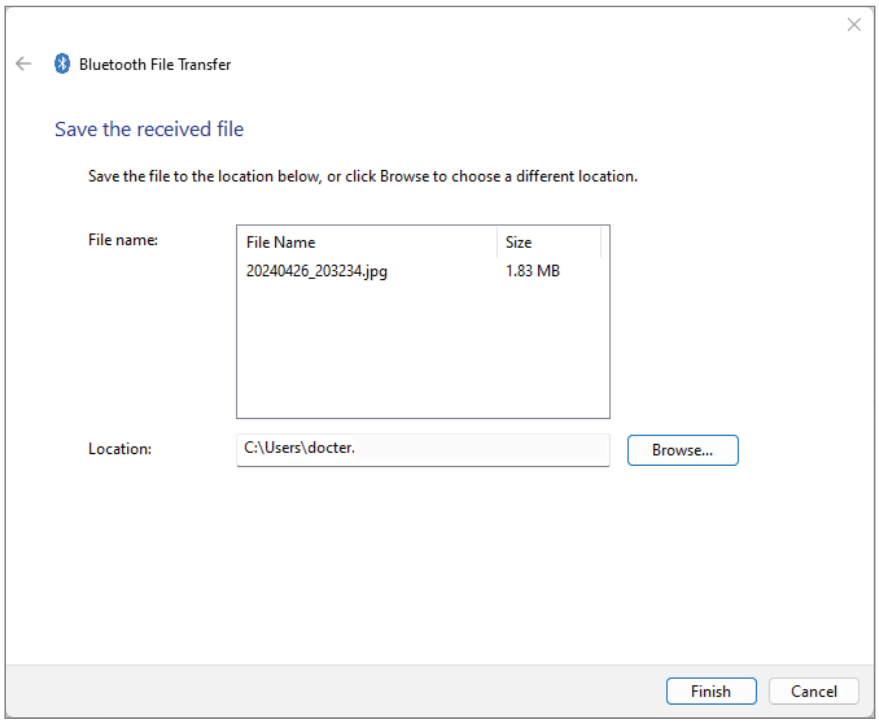

9. Click Finish.

If you plan on doing a lot of file transfers between a mobile device and a laptop or other paired device, it might make sense to get an app to make that job easier. The devices still need to be paired, but the app makes the transfer process easier. Bluetooth File Transfer is one example, and it's available in the Google Play Store. The iOS App Store has several options available as well if you search for Bluetooth file transfer.

Pairing an iPhone

Exercise 10.6 took you through detailed steps on how to pair an Android phone with a Windows laptop and transfer a file. Of course, iOS lets you pair Apple phones with other Bluetooth-enabled devices as well. From the non-phone side, the steps involved are exactly the same.

In iOS, Bluetooth is enabled through Settings ➢ Bluetooth (Figure 10.40). Paired devices are in the My Devices section, and other devices available for pairing will appear under Other Devices. Pairing is done much like it was done from Android OS. Tap the device you want to pair with and begin the pairing process, including confirming the PIN.

FIGURE 10.40 iOS Bluetooth

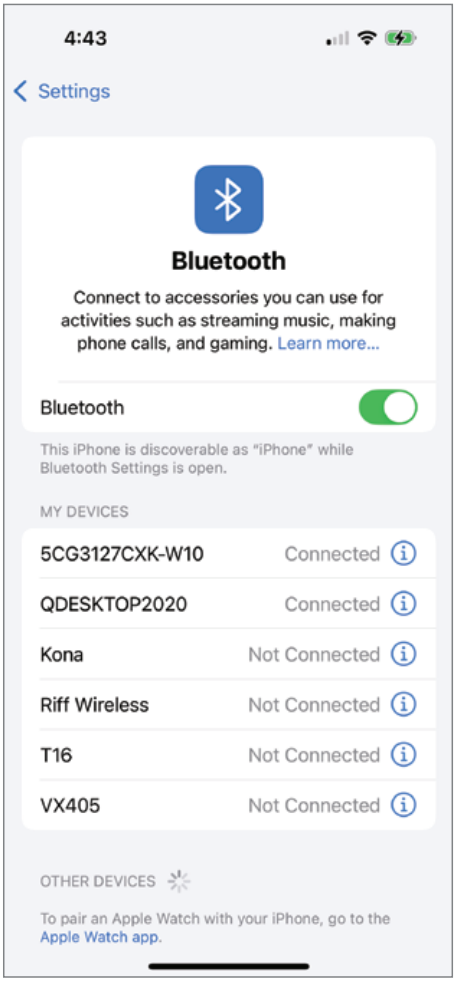
4:43
Settings
Bluetooth
Connect to accessories you can use for activities such as streaming music, making phone calls, and gaming. Learn more...
Bluetooth
This iPhone is discoverable as "iPhone" while Bluetooth Settings is open.
MY DEVICES
5CG3127CXK-W10 Connected
QDESKTOP2020 Connected
Kona Not Connected
Riff Wireless Not Connected
T16 Not Connected
VX405 Not Connected
OTHER DEVICES
To pair an Apple Watch with your iPhone, go to the Apple Watch app.

As far as transferring files goes, it can be easier or more difficult than an Android, depending on the device you're transferring with. If there's another iOS Bluetooth-enabled device within range, it's easy. Open the file or photo or video you want to transfer, tap the share icon (it's a square with an up arrow pointing out of it, in the lower-left corner), and then choose the person's contact with the AirDrop icon or tap the AirDrop icon below (Figure 10.41) and find the person. They will get pinged and must accept the transfer. AirDrop transfers files via Bluetooth Low Energy (BLE).

FIGURE 10.41 Sharing a file in iOS

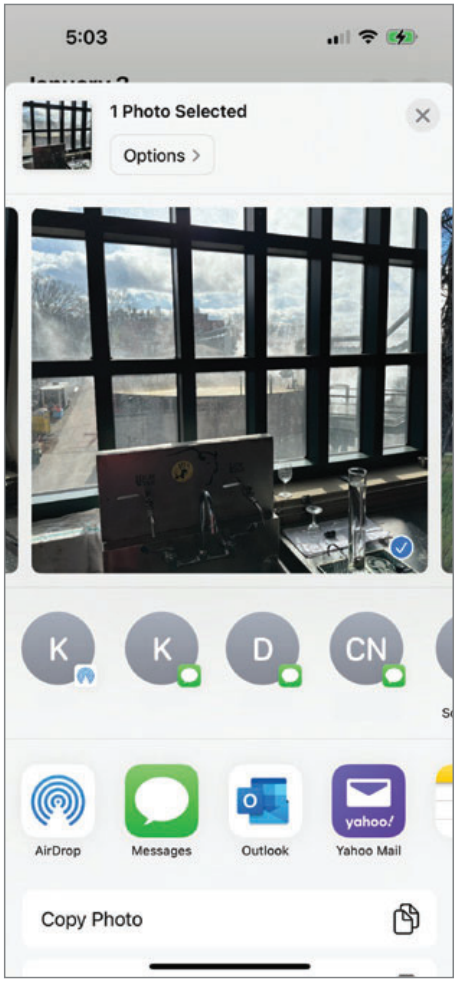

If you want to transfer a file to a non-iOS device, it's more difficult. iOS does not natively support Bluetooth transfers to non-iOS operating systems. There are many third-party apps available to help with this, such as Phone Drive, Air File Share & Drop, and others.

Pairing with simple devices such as headsets, mice, and keyboards that don't have a display is pretty straightforward. We'll use Bluetooth earbuds as an example.

1. Make sure the smartphone (or laptop) is on and Bluetooth is enabled.
2. Turn on the earbuds.

3. You will probably need to press and hold a button on the earbuds to put them into Bluetooth discovery mode. They usually have an LED that will flash a pattern of lights to indicate they are ready for pairing.
4. Once the earbuds appear on the phone's list of available Bluetooth devices, tap on them and confirm pairing.

That's it. The process will be similar for other devices that don't have a display. When in doubt, check the device manufacturer's instructions on pairing.

Understanding Mobile App Support

Mobile devices give users the ability to roam practically anywhere they want to and still be connected to the world. Whether this is a good or bad thing can be up for debate, but here we'll focus on the positive aspects of this freedom. One of the compelling features of mobile devices is they can help pinpoint someone's location and provide directions to somewhere else. This is accomplished through location services, which we will cover in this section.

In many cases, users need to connect mobile devices to corporate networks. This poses a unique set of challenges for network and device management, which requires its own set of principles and policies. We will cover these in a section on mobile device management.

Finally, while mobile devices aren't great for editing spreadsheets or documents, they are more than adequate for sending and receiving email, managing calendars, and storing business and personal contacts. Synchronization services are available to ensure that email, calendars, and contacts are consistent and available for users across their devices. Mobile device synchronization will be the final section of this chapter, where we look at synchronizing to laptop and desktop computers as well as the cloud.

Understanding Location Services

Location services identify where you are and can help give you a route to where you want to be. Two different technologies combine to form what we know as location services, and they are GPS and cellular location services.

GPS and cellular location services are both subobjectives in Exam Objective 1.3. Understand how they work, and be able to enable a mobile device to use location services.

Understanding GPS

Global Positioning System (GPS) is a satellite-based navigation system that provides location and time services. It's great technology for those who are perpetually lost, want to know the best way to get somewhere, or want or need to track down someone else.

The most common commercial use for GPS is navigation; you can get your current location and directions to where you want to go. Other uses include tracking; law enforcement can monitor inmates with location devices, or parents can locate their children via their smartphones. Oil and gas companies use GPS in their geological surveys, and farmers can use GPS-enabled machines to plant crops automatically. Companies also use GPS to track delivery drivers or the location of product shipments. There are three major components to GPS: the satellite constellation, the ground control network, and the receiver. The ground control network monitors satellite health and signal integrity. We'll look at the other two components next.

GPS Satellites

The U.S. Department of Defense (DoD) started developing GPS in the early 1970s, with the goal of creating the best navigation system possible. The first GPS satellite launched in 1978, and today the U.S. government manages 31 total GPS satellites covering the globe. Twenty-four are active satellites for the service, and the rest are backups. Satellites are launched into an orbit of about 11,000 miles above the earth, and old satellites are replaced with new ones when an old one reaches its life expectancy or fails. GPS is free to use for commercial purposes.

There are additional global satellite-based navigation systems managed by other government entities. Collectively, they are called Global Navigation Satellite Systems (GNSSs). All of the systems are outlined in Table 10.1; as you might expect, no two systems are compatible with each other.

TABLE 10.1 Global Navigation Satellite Systems

Name	Managed by	Number of satellites
Global Positioning System (GPS)	United States	31
Global Navigation Satellite System (GLONASS)	Russia	24
Galileo Positioning System	European Space Agency	30
BeiDou Navigation Satellite System (BDS)	China	45
Indian Regional Navigation Satellite System (IRNSS)	India	7

At first glance, it might seem like there are an excessive number of satellites required to run a navigation service. GPS systems were designed to require multiple satellites. Receivers use a process called *triangulation* to calculate the distance between themselves and the satellites (based on the time it takes to receive a signal) to determine their location. They require input from three satellites to provide location or four satellites to provide location and

elevation. Most GNSSs provide two levels of service, one more precise than the other. For example, GPS provides the two levels outlined in Table 10.2.

TABLE 10.2 U.S. GPS service levels

Name	Access	Accuracy	Signal Code	Common uses
Standard Positioning Service (SPS)	Free	3-10 meters, with 95% probability of being accurate within 2 meters	Coarse Acquisition (C/A)	Navigation
Precise Positioning Service (PPS)	Government authorization required, uses encryption	A few centimeters	Precise (P)	Military operations, surveying

The two service levels are separated by transmitting on different frequencies, named L1 and L2. L1 transmits at 1,575.42 MHz, and it contains unencrypted civilian C/A code as well as military P code. L2 (1,227.60 MHz) only transmits encrypted P code, referred to as Y code. In the United States, SPS is free to use; the receiver just needs to manage C/A code. PPS requires special permission from the U.S. DoD as well as special equipment that can receive P and Y code and decrypt Y code. Galileo, in the European Union, provides free open (standard) service, but charges users a fee for the high data throughput commercial (premium) service. Both offer encrypted signals with controlled access for government use.

GPS Receivers

GPS receivers come in all shapes and sizes. Common forms are smartphones (Apple products use the name Location Services), wearable watches and wristbands, stand-alone GPS devices (like the Garmin device shown in Figure 10.42), and ones built into automobiles. Many laptops have built-in GPS capabilities, and you can also find GPS devices that come on a collar for pets. Most stand-alone GPS devices feature capacitive touch screens. The Garmin device shown in Figure 10.42 has a 4.5″ touchscreen; 5″ to 8″ devices are also available. Due to smartphone prevalence and many vehicles having built-in GPS capabilities, stand-alone units such as these have lost popularity.

Understanding Cellular Location Services

Cellular location services is designed to do the same thing GPS does, such as provide a user's location or help navigate a route to a destination. While it uses triangulation just like GPS, the rest of the mechanics are different.

First, while commercial GPS services are free, cellular location services is not. It's provided via subscription from a mobile carrier such as Verizon, T-Mobile, AT&T, and others. Second, instead of using satellites, it uses cell phone towers for its triangulation points. This means that if a user doesn't have cell phone reception, then cellular location services won't work.

Cellular location services is also less precise than GPS. Recall that GPS is accurate within about 10 meters, but cellular location services may only be accurate within 100 meters, depending on the user's distance from the cell tower and signal strength. If a user's phone is within range of multiple cell towers, that precision can increase substantially and get close to GPS performance. However, the rule of thumb is that for outdoors navigation, GPS is more accurate than cellular when locating devices.

FIGURE 10.42 Garmin Nuvi GPS

While outdoors, mobile devices can use both GPS and cellular for location services. They might even use a combination of the two, provided of course the device can find satellites or cell towers.

When in range (usually indoors), mobile devices can use Wi-Fi and Bluetooth for location services as well. Both Wi-Fi and Bluetooth can be extremely accurate when pinpointing devices.

Configuring Location Services

Knowing how to turn location services on or off on a GPS receiver or a mobile device is a valuable skill. The specific settings depend on the operating system, of course, but we can provide general guidance here. Exercise 10.7 shows you how to configure Location Services in iOS 18.

EXERCISE 10.7

Configuring Location Services in iOS

1. Click Settings ➢ Privacy & Security ➢ Location Services. You will see a screen similar to the one shown in Figure 10.43.

FIGURE 10.43 iOS location services

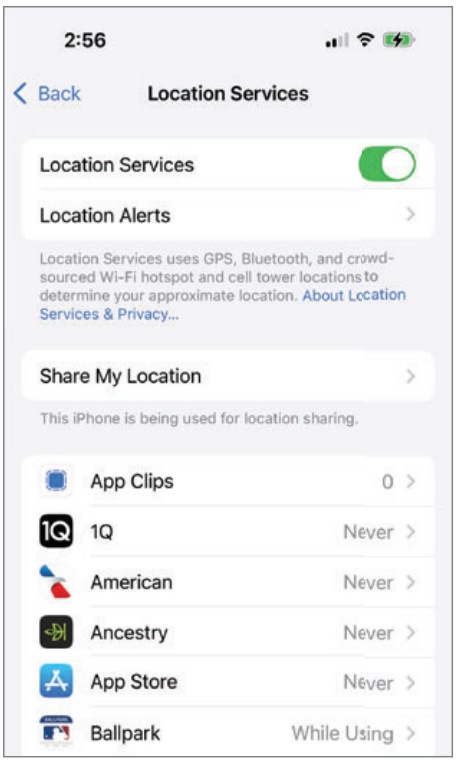
2:56
Back
Location Services
Location Services
Location Alerts
Location Services uses GPS, Bluetooth, and crowd-sourced Wi-Fi hotspot and cell tower locations to determine your approximate location. About Location Services & Privacy...
Share My Location
This iPhone is being used for location sharing.
App Clips 0
1Q Never
American Never
Ancestry Never
App Store Never
Ballpark While Using

2. At the top, there is a slider to enable Location Services. Slide it to the on position to enable it (like in Figure 10.43), and off to disable it.

Note that if you disable it, none of the apps on the device will be able to use the feature, except for Find My iPhone. When on, each app's capability to use Location Services is configured individually.

3. To change an individual app, tap on it. There will be four choices:
   - Never.
   - Ask Next Time Or When I Share.
   - While Using the App.
   - Always.
4. Configure each app as desired.

In Android OS, location services is configured through Settings as well. Exercise 10.8 walks you through how to configure it.

EXERCISE 10.8

Configuring Location Services in Android

1. Tap Settings ➢ Location. You will see a screen similar to the one shown in Figure 10.44.

FIGURE 10.44 Android location settings

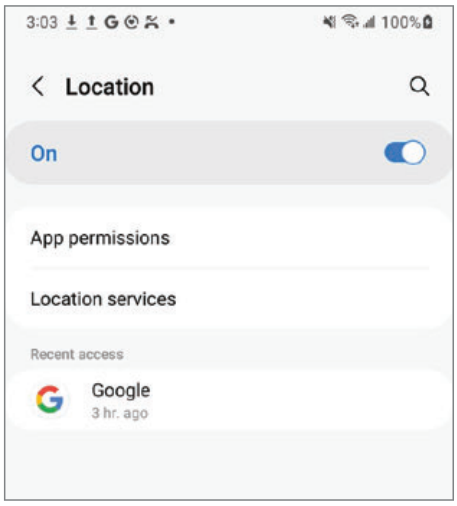

2. Move the slider to the on position, as shown, to enable location usage.

   You can also configure things such as individual app permissions, and the ability to use Bluetooth and Wi-Fi to improve accuracy (by tapping Improve Accuracy), and location services for emergency purposes.

EXERCISE 10.8 *(continued)*

3. Tap App permissions to set location services for each app (Figure 10.45). Options include:
   - Allow All The Time.
   - Allow Only While Using The App.
   - Ask Every Time.
   - Don't Allow.
4. On the Location screen, tap Location Services to configure emergency services, location details, and the use of Wi-Fi and Bluetooth for location purposes (also Figure 10.45).

FIGURE 10.45 Android location (a) app permissions and (b) location services settings

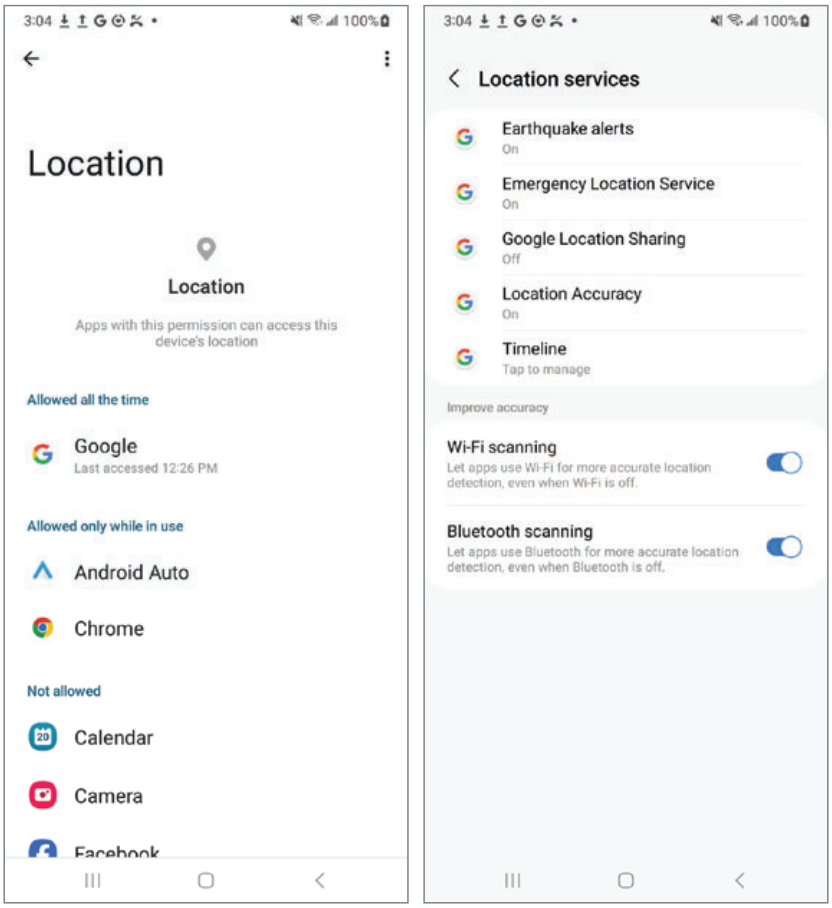

5. Configure apps and location options as you would like, and close Settings.

Understanding Mobile Device and App Management

The use of mobile devices on corporate networks has increased exponentially over the past several years, and it will likely continue to increase over time. So many people work remotely or travel that it just makes sense for them to be able to manage email or perform tasks such as order placement and management from their mobile devices. Of course, these usage cases can give network administrators nightmares, as every additional device connected to the network represents another security risk. And when devices are small and easily misplaced or stolen, that elevates the security risk to a whole new level.

To help with mobile device security, many companies use a combination of mobile device management and mobile application management. We'll look at those in the upcoming sections, including device configurations, policy enforcement, and app management. We will finish this section off with a deep dive into the most common corporate use of mobile devices, which is sending and receiving email.

Mobile Device Management and Mobile Application Management

Imagine that you are a network administrator for a corporate network, and the company implements a new policy where mobile devices should be granted network access. As we mentioned earlier, if done incorrectly this can pose a massive security risk to the company, so, no pressure, right? With security in mind, you may want to explore implementing a *mobile device management (MDM)* solution.

Exam Objective 1.3 lists several topics to know for mobile device management. They include corporate versus BYOD device configurations, policy enforcement, and corporate applications.

An MDM is a software package residing on a server. The key purpose of an MDM is to enroll mobile devices on the corporate network, and once those devices are enrolled, to manage security. This is done through security policies as well as the ability to remotely track, lock, unlock, encrypt, and wipe mobile devices as needed. Now if someone's smartphone is misplaced or stolen, an administrator can wipe it remotely and the security threat is mitigated.

Although this is a good solution for device-level security, there's a big piece missing—the software. That's where *mobile application management (MAM)* comes into play. Typically implemented in conjunction with an MDM, an MAM allows network administrators to remotely install, delete, encrypt, and wipe *corporate applications* (apps developed by the company and/or approved for use on company devices) and related data from mobile devices. In an MAM, administrators can specify software packages that are allowed to be installed on the mobile device and prohibit others that could pose security risks. When enrolled in the MAM, users may have a corporate app store that functions similarly to Apple's App Store or Google Play Store. Figure 10.46 shows an example of some apps in a corporate app store, managed by VMware Airwatch.

FIGURE 10.46 Corporate app store

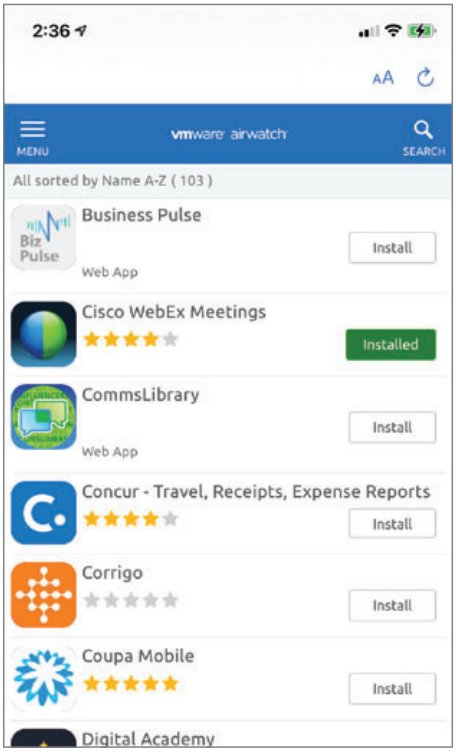

Device Configurations

On a corporate network, there are two classifications of mobile devices: corporate devices and bring your own device (BYOD). A *corporate device* is one owned by the company and given to the user for business use. In many cases, the device is strictly limited to business use, or at least it's supposed to be. The only apps allowed on the device will be corporate approved apps. The user might not even be able to take personal photos or videos with the device.

To reduce costs needed to provide employees with mobile devices, many companies have adopted a *bring your own device (BYOD)* plan. Under the plan, employees can use their personal smartphones, as long as the phone meets certain security requirements, and install corporate applications. In situations like these, the MDM will be responsible for granting the phone access to the network. If the phone gets lost, however, it may be impractical to use an MDM to wipe it because it would erase all of the user's personal files (photos, music, apps, etc.) as well. With the use of an MAM, though, the corporate applications and data on the lost phone can be remotely wiped, eliminating the security risk for the company.

MDM Policy Enforcement

Any organization that allows mobile devices on its network should have an MDM acceptable use policy (AUP) in place. Users must be trained on the policy, understand the

consequences of not following it, and acknowledge their awareness of it in writing or electronically. Topics in an MDM AUP typically include:

- Use of strong authentication and encryption methods
- The list of approved apps (or where to get approved apps)
- Policies for corporate data storage, including personally identifiable information (PII)
- Policies for handling lost or stolen devices, including data wipes
- Acceptable use of the device in public places, including the use of public Wi-Fi

The MDM software management package will have tools available to monitor corporate and BYOD devices for compliance. If a device is found to be out of compliance, the MDM can automatically fix the issue by locking the device or removing the offending material without interaction from the user. It can also send a notification to the administrator, who can then handle the situation per company policy. For example, the user's manager may also receive the notification and then take appropriate action.

Although an MDM can automatically wipe a BYOD device, it's usually frowned upon because the user may have unrecoverable personal data on that device. Instead, the MDM can selectively wipe data from the device, leaving the user's personal information intact. Completely wiping a BYOD device should only be done as a last resort.

Configuring Email Accounts

Accessing email is the most common corporate use for mobile devices. Usually, the most difficult part is finding the server settings that are used only during establishment of the connection, which tends to occur only once for each device. Another big challenge is when users have the same devices and accounts for many years; since they have not configured them in some time, it might be difficult for them to remember their usernames and passwords, if needed.

When configuring mobile devices to access email, you will be attaching to one of the following two types of services:

- An integrated commercial provider, which includes Gmail, iCloud, Exchange Online, Yahoo Mail, and others
- A corporate or ISP-based email service

Connecting to an integrated commercial provider is generally quite easy and nearly automatic. Connecting to a corporate or ISP-based account usually involves a few more steps, but it shouldn't be too tricky if you know the proper server settings. In the following sections, we'll look at configuring email on a mobile device and settings to know for manual email configuration.

Mobile Internet Email Configuration Options

If your email is on a common web-based service, such as Gmail, iCloud, Exchange Online (Microsoft Exchange), or Yahoo, configuring the email feature is pretty straightforward. Usually, your email address and password are all that are required. However, if you have a corporate or ISP account or a custom domain, even if it's hosted and accessible through Gmail, `Outlook.com`, or the other popular services, you may need to take a few more steps to make a connection. For the purposes of trying out a commercial service, you can always make a dummy account on the service's website and play around with that if it helps you complete these exercises. Exercise 10.9 and Exercise 10.10 detail the basic steps required to configure a commercial email account on an iPhone and on an Android standard email client, respectively.

EXERCISE 10.9

Email Account Configuration on an iPhone

1. From the home screen, tap Settings ➢ Apps ➢ Mail and then tap Mail Accounts. You will see a screen similar to the one shown in Figure 10.47.

FIGURE 10.47 Email accounts

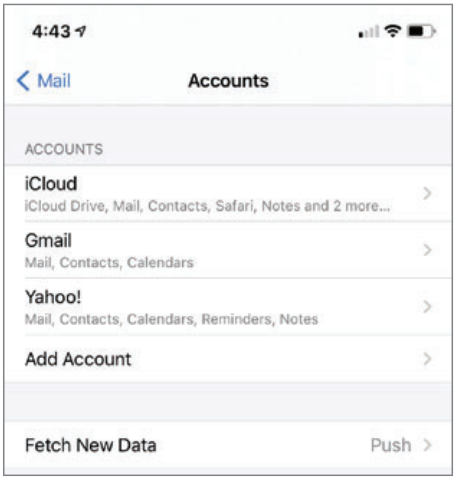

   If the device already has configured accounts, they will be listed under Accounts. For example, this device already has iCloud, Gmail, and Yahoo email accounts on it.

2. Tap Add Account. A screen similar to the one shown in Figure 10.48 will appear.

FIGURE 10.48 Adding an account

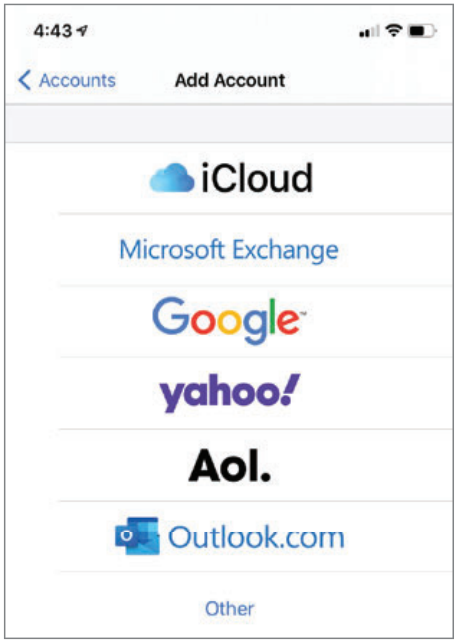

Note that if you have an email account with any of the listed services, you can tap it to configure your client. You will be asked for your username and password. After entering this information, you should be done. It's as easy as that to add an account for a commercial provider.

If the account is with a different provider, a few more steps are needed. They are as follows:

3. Tap Other ➢ Add Mail Account.

   You will get a screen similar to the one shown in Figure 10.49.

FIGURE 10.49 Creating a new email account

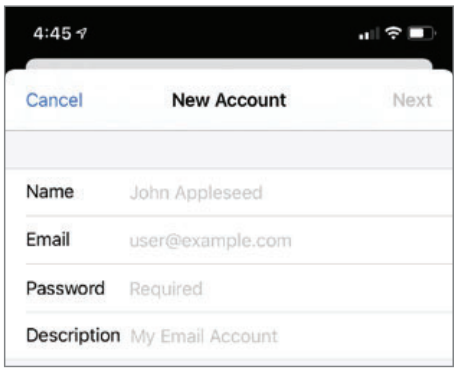

EXERCISE 10.9 *(continued)*

4. Add your name, email address, password, and an optional description.
5. Tap Next.
6. In the New Account screen (see Figure 10.50), choose IMAP or POP.

   You were introduced to IMAP and POP in Volume 1 Chapter 7, and we will review the differences between the two in the next section. If your server supports it, IMAP is preferable.

   Here you also configure the names of the incoming and outgoing mail servers. Many organizations will have one server handle both functions.

FIGURE 10.50 New Account screen

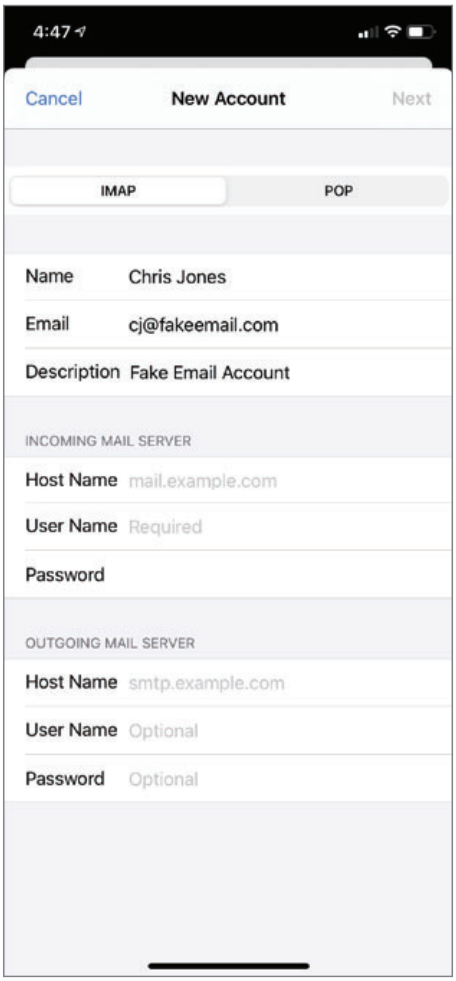

7. Once you have entered the correct information, tap Next.

   The iPhone will make connections to the server(s) to verify the username and password. If the username and password are successful, you will get a screen asking which content you want to receive, such as mail, contacts, and calendars.

8. Tap Save, and it will add your account.

   Now you can use the Mail icon on the home page to retrieve email. Note that you can have several email accounts configured to receive email in this app (as was the case in Figure 10.47).

9. (Optional) To delete an account, repeat steps 1 and 2 in this exercise. On the Passwords & Accounts screen, tap the account that you want to delete.

10. (Optional) Tap Delete Account to remove it from this client.

Exercise 10.10 details the steps required for configuring an email account on an Android standard email client. If the Android device does not have the email app on the home screen, you can add it or run it directly from the All Apps list.

EXERCISE 10.10

Email Account Configuration in Android

1. Tap Settings ➢ Accounts And Backup ➢ Manage Accounts (see Figure 10.51).

FIGURE 10.51 Accounts in Android

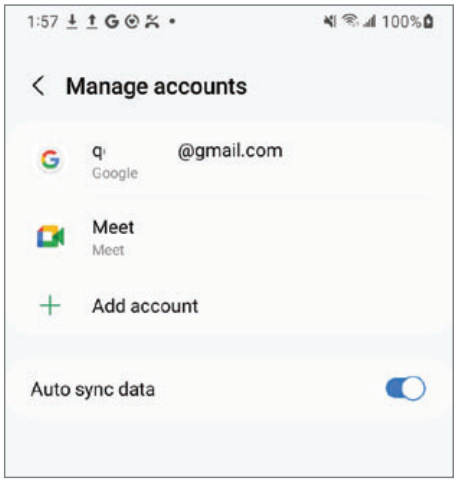

2. Tap Add Account. The list will look similar to the one shown in Figure 10.52.

EXERCISE 10.10 *(continued)*

FIGURE 10.52 Adding an account

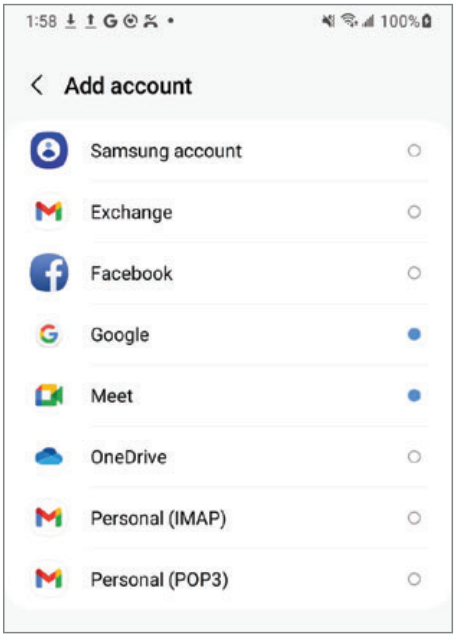

3. Tap the type of account that you want to add. For this example, choose Personal (IMAP).
4. Enter an email account name and tap Next. (For purposes of this exercise, you can make up fake credentials—the rest of the exercise will work through step 9 with them.)
5. Enter a password and tap Next.

   You will get a screen similar to the one shown in Figure 10.53.

FIGURE 10.53 Incoming Server Settings

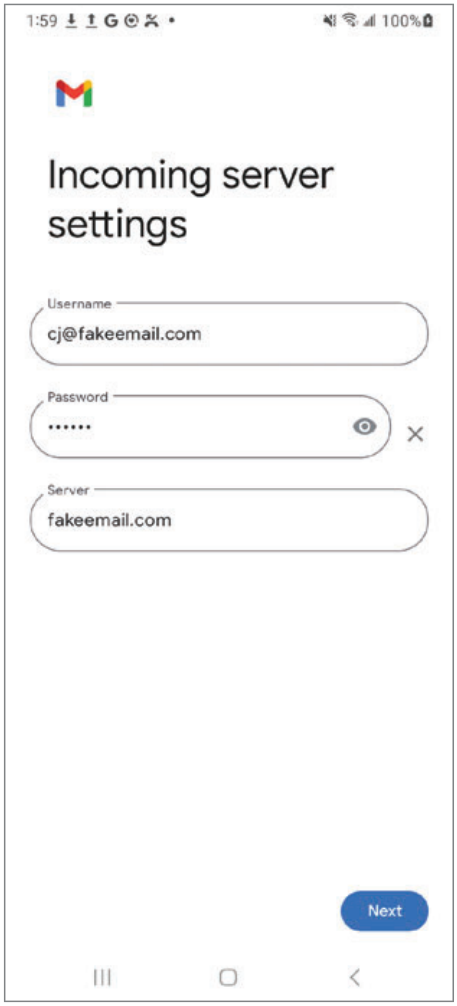

6. Tap Next to validate the server settings. It may take a few minutes to complete.

   The information we entered is obviously fake, so the validation will fail. (Hopefully yours succeeds if you used real credentials!) You will get a warning saying that email security is not guaranteed. Click Edit Settings. Having the validation fail gives us a chance to talk about additional options.

   Since the validation failed, Android tries to help out by suggesting we try updating the port or security type (see Figure 10.54). Notice that near the bottom of the screen, the port is 143 because we chose IMAP in step 3. Continue with the following steps:

EXERCISE 10.10 *(continued)*

FIGURE 10.54 Email port setting

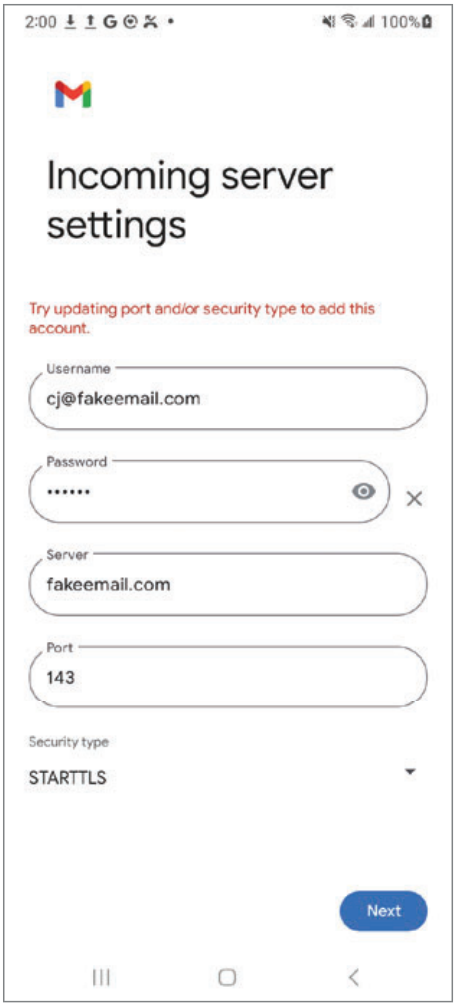

7. Tap anywhere on the setting under Security Type, and you can change the security to SSL/TLS.

   Options are shown in Figure 10.55. Changing the type to SSL/TLS will automatically change the port number to 993. (We will discuss SSL/TLS shortly.)

FIGURE 10.55 Email security options

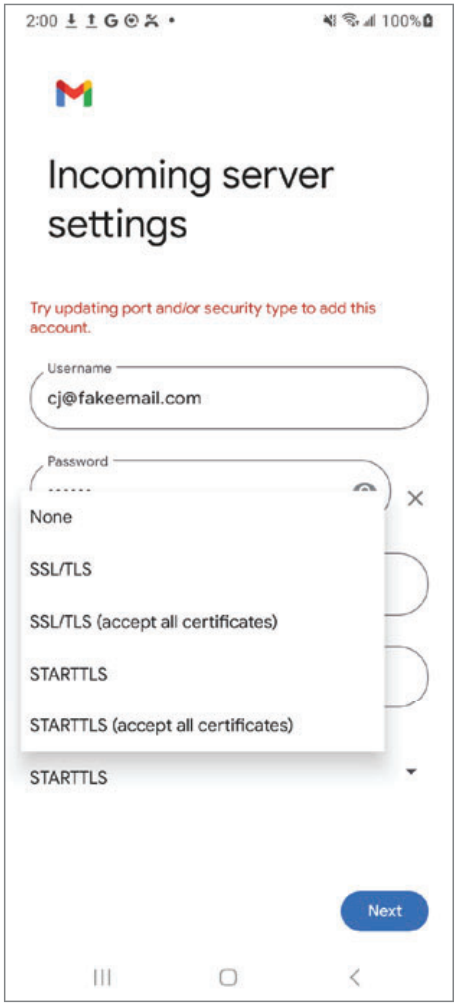

8. Tap Next. Android will contact the server and verify the server settings.

   After verifying a connection to the server, you may be able to choose account options, such as sync frequency and email notification settings.

9. Choose your settings (if applicable) and tap Next.

   You may get an account confirmation page, where you can edit the account name that gets displayed as well as the name that gets displayed on outgoing messages.

10. Tap Next.

    The new account will appear in your accounts list in Settings.

11. Access your email through the Email app on the home page.

Manual Email Configuration Options

In situations when you find that your email client cannot automatically configure your email account for you, there are often manual settings for the protocols required for sending and receiving emails. Table 10.3 outlines these protocols and their uses. These should look familiar to you if you recall Volume 1 Chapter 7.

TABLE 10.3 TCP/IP mail protocols

Mail protocol	Description	Default port number
Simple Mail Transfer Protocol (SMTP)	Used to communicate between client and server and between servers to send mail to a recipient's account. The key word is *send*, as this is a push protocol.	TCP 25
Post Office Protocol (POP)	Used to communicate between a client and the client's mail server to retrieve mail with little interaction.	TCP 110
Internet Message Access Protocol (IMAP)	Used to communicate between a client and the client's mail server to retrieve mail with extensive interaction.	TCP 143

In Table 10.3 and for the rest of this section, we refer to the POP and IMAP protocols. In practice, the current versions of each are POP3 and IMAPv4. For brevity, we will refer to them as just POP and IMAP—the version numbers do not affect the port numbers and relevant configuration information.

In a TCP/IP network using only the protocols in Table 10.3 (there are other, less common options), you must always use SMTP for sending mail. You must decide between the use of POP and IMAP for interacting with the mail server to retrieve your mail with the client. When supported, IMAP is a clear choice because of its extensive interaction with the server, allowing the client to change the state or location of a mail item on the server without the need to download and delete it from the server.

Conversely, POP limits client interaction with the server to downloading and deleting items from the server, not allowing their state to be changed by the client. In fact, the use of POP as your receive-mail protocol can lead to confusion because copies of the same items appear in multiple client locations, some marked as read and others unread. Additionally, where IMAP changes the state of a mail item on the server and leaves the item on the server for later access by the same or a different client, POP settings must be configured *not* to delete the item from the server on download to each client. This, however, is what leads to the choice of multiple copies among the clients or only one client being able to download the item.

Most, if not all, Internet mail services require secure connections. SMTP, POP, and IMAP are all unsecure protocols, so this poses a problem. One solution is to use *Secure Sockets Layer (SSL)* or *Transport Layer Security (TLS)* on top of these protocols. You might recognize SSL and TLS from their use on TCP port 443 for the HTTPS protocol. When using these protocols to secure the email protocols, mail servers and clients need to communicate on ports other than the ones shown in Table 10.3. These are outlined in Table 10.4.

TABLE 10.4 Secure mail ports

Mail protocol	TCP port number
SMTP with SSL	465
SMTP with TLS	587
IMAP with SSL/TLS	993
POP with SSL/TLS	995

Using STARTTLS

It may be mildly obnoxious to have to memorize multiple port numbers (although you should for the exam), as well as program email clients and servers to recognize different ones. A different solution exists, called *STARTTLS*. (You saw it as an option in Figure 10.55.)

STARTTLS secures email communications over the three common email protocols, but it does not change their port numbers. And despite its name, STARTTLS doesn't exclusively use TLS—it can use SSL as well. Some email software supports it, whereas others do not, so be aware of its existence and know what it does. It's not a test objective, but it's something you may run into in the real world.

Additionally, you will need to know the server names for your service. Sometimes they are the same for inbound and outbound mail handling, but they may be different. Table 10.5 lists the servers in the United States for iCloud, Gmail, Exchange Online, and Yahoo Mail. Unless otherwise specified, the ports in Table 10.4 should be used for the protocols listed in Table 10.5.

TABLE 10.5 Secure mail servers for common email services

Service	Direction and protocol	Server name
Gmail	Outbound on SMTP with SSL or TLS	`smtp.gmail.com`
	Inbound on IMAP with SSL	`imap.gmail.com`
	Inbound on POP with SSL	`pop.gmail.com`
iCloud	Outbound on SMTP with SSL	`smtp.mail.me.com`
	Inbound on IMAP with SSL	`imap.mail.me.com`
Exchange Online	Outbound on SMTP with TLS	`smtp.office365.com`
	Inbound on IMAP with SSL	`outlook.office365.com`
	Inbound on POP with TLS	`outlook.office365.com`
Yahoo Mail	Outbound on SMTP with SSL	`smtp.mail.yahoo.com`
	Inbound on IMAP with SSL	`imap.mail.yahoo.com`
	Inbound on POP with SSL	`pop.mail.yahoo.com`

Configuring Mobile Device Synchronization

While some users are content to have a mobile device only, many others do not consider their mobile devices to be islands unto themselves. Instead, they treat their mobile devices as extensions of their primary computing devices that, even if they happen to be portable, stay at work or at home while the mobile devices go on the road with the users. However, because many of the same changes to a user's calendar, contacts, and personal files can be made from the mobile device as easily as from the primary computer, frequent synchronization of the two devices is in order. *Synchronization* is the act of mirroring all unique changes and additions from one device to the other.

In most cases, there are multiple options as to how the mobile device will connect to the computer system it is syncing to. Common connections include over USB, across Wi-Fi, over Bluetooth, and through a cellular connection. Although the wired connections tend to be the most reliable, the convenience of wireless connections and their automatic unattended synchronizations cannot be ignored. When syncing, you can either sync to the cloud or a local computer (such as a laptop, desktop, or networked server).

Because each manufacturer of mobile devices must approach synchronization of data in the best manner for their devices, generalized discussions of data to be synchronized can *only* include the common types. These include items such as the following:

- Calendars
- Contacts
- Business applications (email and productivity apps)
- Cloud storage

In the following sections, we will look at each of these. We will also provide guidance on how to sync iOS and Android devices and understand data caps.

Recognizing Data Caps

Syncing a mobile device to the cloud or a local computer is smart mobile device management. Mobile devices can easily be misplaced or broken, yet many of them have apps and data that would be hard to replace, or photos that would be impossible to replace. Where you choose to sync to is a matter of personal preference.

Recognizing data caps is part of Exam Objective 1.3. Understand what they are and how to overcome them.

If you sync to the cloud or use a cellular connection to sync, you can run into a problem with *data caps*, or limitations to the amount of data you can transfer. The sync itself might not consume too much data if you have a limited data plan, but when added to email downloads, streamed videos and music, and GPS use, the caps may prove to be too restricting. Here are some tips to see if you have data caps and if so, things you can do about it:

- Research your contracts with your cellular provider and ISP to see if there are data caps, and if so, if your monthly usage comes close to getting capped out.
- Most ISPs have an app you can use to track usage. On mobile devices, iOS device users can go to Settings ➢ Cellular to see how much data has been used, and Android device users can do the same under Settings ➢ Connections ➢ Data Usage ➢ Mobile Data Usage.
- Research current and other providers to see if they have caps, and the cost associated with plans that don't have caps.
- Change settings to help preserve data. For example, do you need to stream videos in 4K on your tablet? Lowering the video quality will consume less data, and most streaming services have options to do so.
- Take advantage of off-peak hours. Some providers allow for extra data during early morning hours when fewer people are on their devices. If you can, that's a perfect time to schedule a major download or data sync.

With a little legwork and planning, you can avoid data caps and any negative effects they might have on your mobile device experience. It's probably easiest to spend a little more and upgrade your plan to an unlimited one if it's available.

Syncing Apple iOS Devices

Apple provides specific tools for its iOS and iPadOS users to sync and backup their devices to the cloud and to a desktop or laptop computer. The differences between syncing and backing up aren't many. Synchronization means that the exact same copy of the data (music, pictures, contacts, or whatever) is on both the iOS device and the local computer. Backing up means taking whatever is on the phone at that time and ensuring that a duplicate is stored elsewhere. Synchronization can often happen both ways, whereas backups are a one-way process. The two tools Apple provides are iCloud and iTunes.

Using iCloud

Apple's version of the cloud is called iCloud, and it is available to all iOS users. We will cover details on what the cloud really is in Volume 1 Chapter 11, "Virtualization and Cloud Computing"; for now, consider it to be storage space accessed through the Internet. iCloud is the only sync option for users who have ditched their desktops, and a convenient option for those who haven't.

When a user creates an Apple ID, it's used to log into iCloud (and iTunes, which we will discuss in the next section). Apple recommends that the same username be used for both, but it is not required. On the iOS device, it's easy to get to iCloud settings. Open the Settings app, tap the Apple ID at the top of the page, and then choose iCloud. The configuration page is shown in Figure 10.56.

At the top of the iCloud settings, it will show you the space available for that Apple ID. The default amount of free storage space is 5 GB. Tapping Upgrade to iCloud+ will give you an option to purchase more space if needed, such as 50 GB, 200 GB, or 2 TB.

In the Saved to iCloud section it shows which types of data are being synced. To change the setting for an app, tap it (such as the Photos app shown in Figure 10.57) and then change the toggle for Sync this iPhone. If the app you want isn't shown on this page, tap See All and it will give you an extensive list. To back up the phone, tap iCloud Backup and set the configuration options you want.

Synchronization and backups will happen when the phone is plugged into a power source, locked, and connected to Wi-Fi.

Using iTunes

To sync a device with a desktop or laptop, you must have the iTunes app installed on your computer. It's installed by default on Macs and can be found at `www.apple.com/itunes` for non-Apple OSs. Figure 10.58 shows the Summary page of iTunes for an iOS device when it's connected. Notice that the Backups section has options to back up to iCloud or to This Computer. In this section, we will focus on local backups.

FIGURE 10.56 iCloud configuration settings

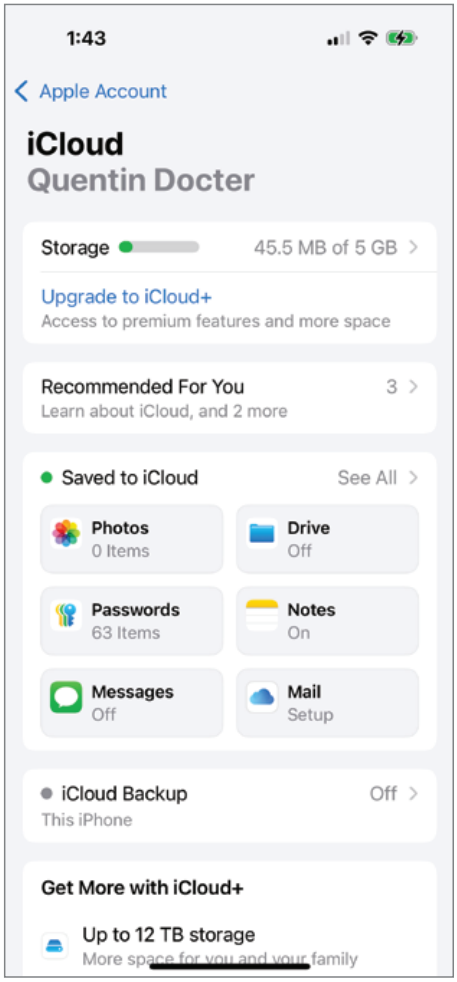

By default, iOS devices will automatically sync each time they are connected by USB, and Wi-Fi if it's been enabled, and they are recognized under the Devices section in the left frame of iTunes. The exception is when iTunes is set to prevent automatic synchronization. Figure 10.59 shows the dialog box from iTunes attained by clicking Edit ➢ Preferences ➢ Devices. (You might need to press the Alt key to get the Edit menu to appear.) Notice that syncing is set to occur automatically because the Prevent iPods, iPhones, And iPads From Syncing Automatically check box is cleared.

The selection of what is to be synchronized is a task unto itself, but iTunes provides specific tabs on the left side of the interface for each class of data, as shown in Figure 10.58, under the Settings section.

FIGURE 10.57 iCloud Photos settings

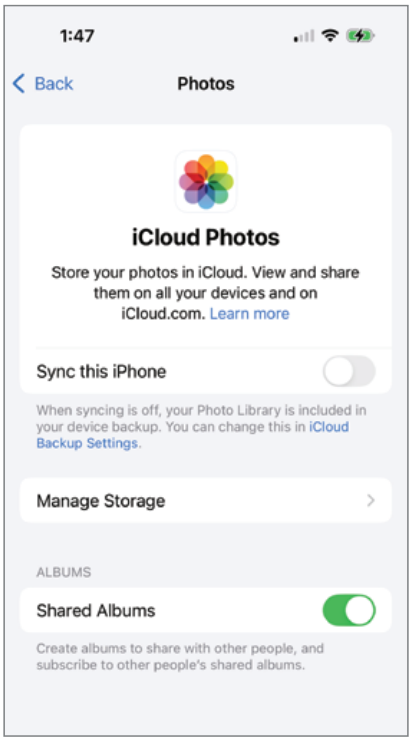
1:47
Back
Photos
iCloud Photos
Store your photos in iCloud. View and share them on all your devices and on iCloud.com. Learn more
Sync this iPhone
When syncing is off, your Photo Library is included in your device backup. You can change this in iCloud Backup Settings.
Manage Storage
ALBUMS
Shared Albums
Create albums to share with other people, and subscribe to other people's shared albums.

FIGURE 10.58 iTunes Summary page

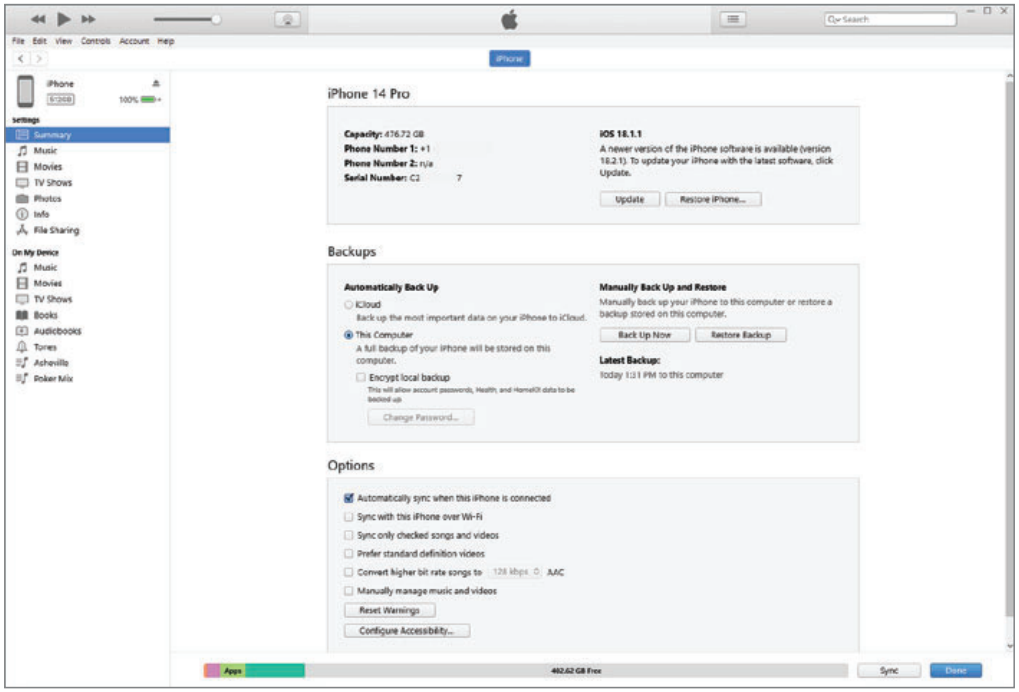

FIGURE 10.59 Devices Preferences in iTunes

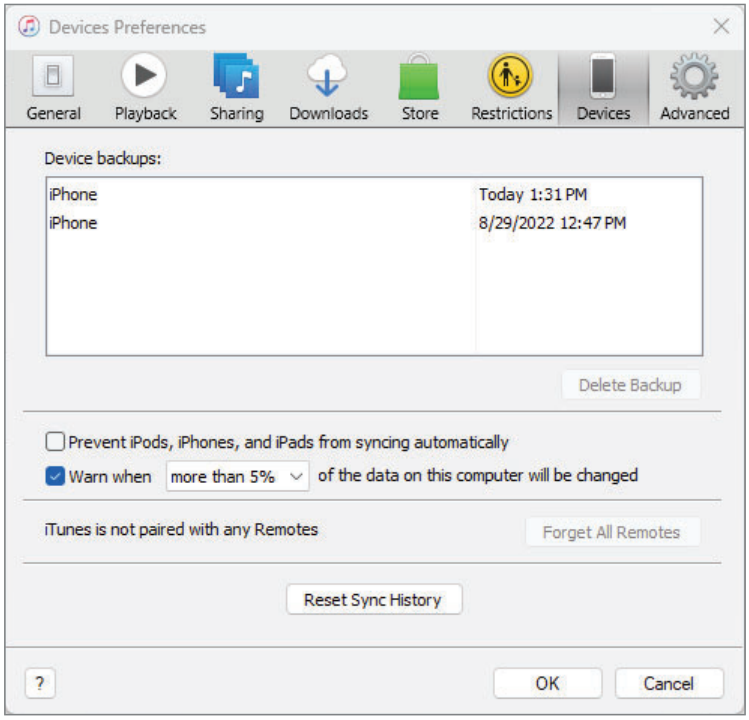

If automatic synchronization is disabled in Devices Preferences, you can start the manual synchronization of an iOS device by selecting the iOS device above the left frame in iTunes and then clicking the Sync button at the bottom-right corner of its Summary tab.

Syncing Android Devices

Just as with Apple's devices, mobile devices built for the Android operating system can be synced to the cloud or a traditional computer. Apple's iTunes is proprietary and has been designated as the application that performs synchronization of iOS devices. In a similar way, manufacturers of Android devices have their own syncing utilities. Because this software and the connection methods allowed vary widely from one manufacturer to another, it is difficult to predict exactly what one manufacturer will offer in its utility and whether each Android device it produces will interact the same way and over the same connections.

Android devices often use Google Drive as the preferred cloud storage location. Manufacturers will have their own backup and sync tools, such as Samsung SideSync, LG PC Suite (also known as LG PC Sync), and HTC Sync Manager.

Let's use a Samsung phone as an example. If you want to configure backups using Google Drive, tap Settings ➢ Accounts And Backup ➢ Back Up Data (Figure 10.60). Notice on this phone that there is a Samsung Cloud section as well as a Google Drive section. It's user preference as to which cloud you choose, but for this example we will use Google Drive. Tapping Back Up Data under Google Drive takes us to a screen like the one shown in Figure 10.61. In this instance, backup is enabled and the device is automatically backed up. To run a manual backup, tap the Back Up Now button.

FIGURE 10.60 Samsung Accounts and Backup

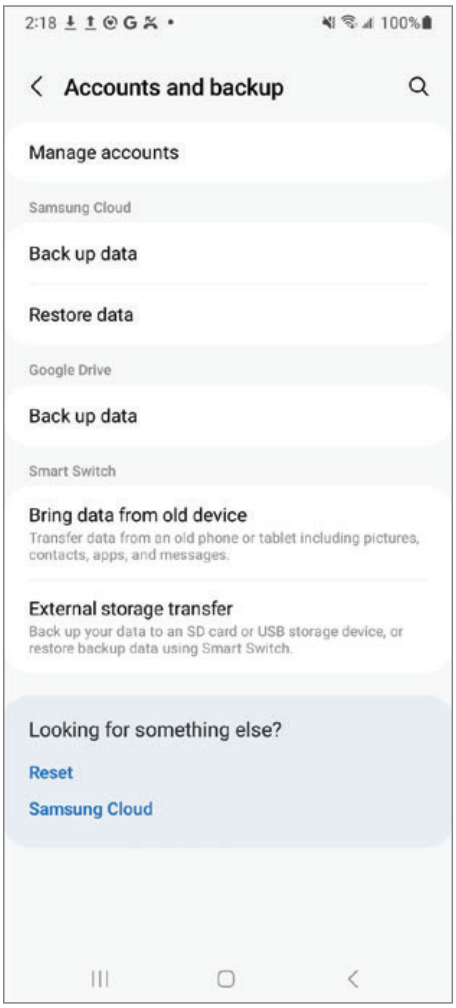
2:18
100%
Accounts and backup
Manage accounts
Samsung Cloud
Back up data
Restore data
Google Drive
Back up data
Smart Switch
Bring data from old device
Transfer data from an old phone or tablet including pictures, contacts, apps, and messages.
External storage transfer
Back up your data to an SD card or USB storage device, or restore backup data using Smart Switch.
Looking for something else?
Reset
Samsung Cloud

Common items available for synchronization include contacts, applications, email, pictures, music, videos, calendars, bookmarks, documents, location data, social media data, e-books, and passwords. Of course, specific options will depend on the software used.

FIGURE 10.61 Google Drive backup

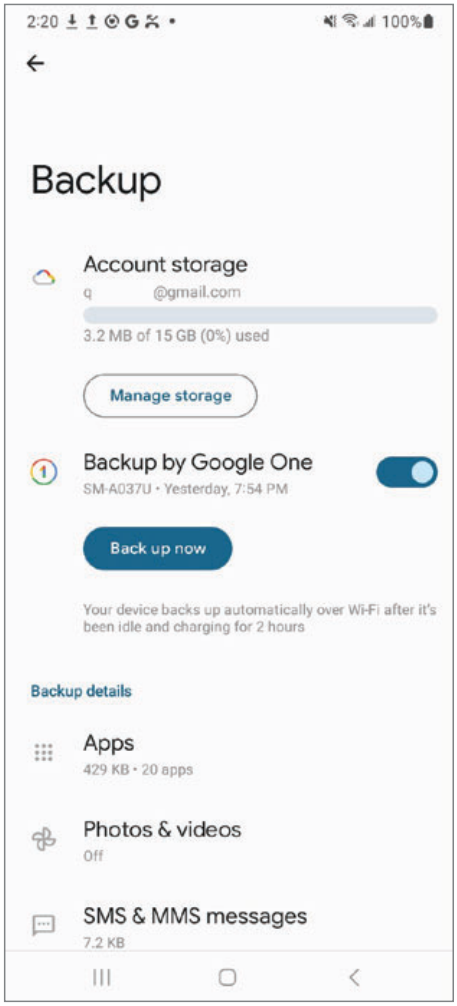

Syncing Calendars, Contacts, and Business Applications

For the most part, users who follow the syncing advice for their iOS or Android device will be able to sync or back up anything they need to. There are a few cases where that might not be true, however. For example, email, calendars, and contacts might require different sync settings. In addition, some business apps like Microsoft 365 have their own sync settings as well.

Exam Objective 1.3 includes knowing how to synchronize calendars, contacts, and business applications such as email. Be familiar with these concepts for the exam!

Email, Calendars and Contacts

Most users will use their email app to manage calendars and contacts. As long as email is set up to synchronize between the server and the mobile device, calendars and contacts will automatically sync as well. Those features may be adjusted and synced individually within the email app. Fortunately, all major commercial email providers sync between their servers and users' devices automatically. This includes users who have multiple devices to access the same email accounts. Everything is stored on the server, so it shouldn't matter which device the user is on. Of course, there are some exceptions to the rule. For example, if a user is writing an email on their smartphone and it loses power or its network connection before the email autosaves, the draft email might not be available on the server (or other devices) until connectivity is restored.

Regardless of the service being used, synchronization is always based on the user's account ID. For example, a Gmail user will use their `@gmail.com` address, and an iCloud user will use their `@icloud.com` address.

In iOS, the synchronization settings for an email account can be checked by going to Settings ➢ Apps ➢ Mail and then tapping the account. You will see a screen like the one shown in Figure 10.62.

FIGURE 10.62 iOS Sync data types

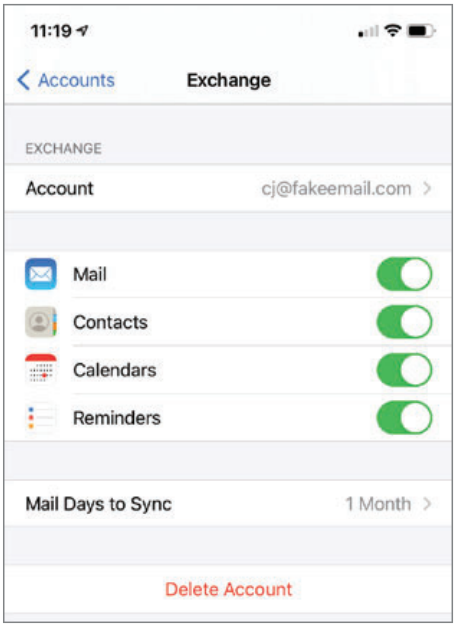

In Android, the same settings are found in Settings ➢ Accounts and Backup ➢ Manage Accounts. There you will see the list of accounts (it was shown in Figure 10.51) on the

device. Tap the account name (the Google account, in this case), and then tap Sync account. That shows what items are being synced, as in Figure 10.63.

FIGURE 10.63 Android Sync Account

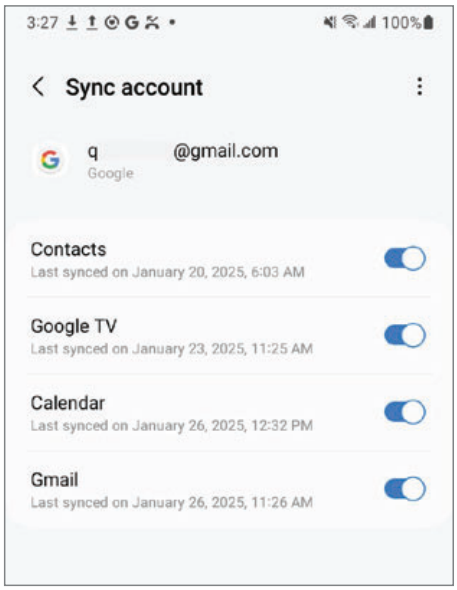

In some cases, iOS users will want to sync the Calendar or Contacts app on their phones with their mail service. To do this for the Calendar, go to Settings ➢ Apps ➢ Calendar and tap Calendar Accounts (Figure 10.64). Then tap the name of an account to sync with it, or tap Add Account to add a new account to sync to. If you tap Add Account, the process will be just like adding an email account like you did in Exercise 10.9. Contacts works the same way—just choose the Contacts app instead of Calendar app in Settings ➢ Apps.

FIGURE 10.64 iOS Calendar Accounts

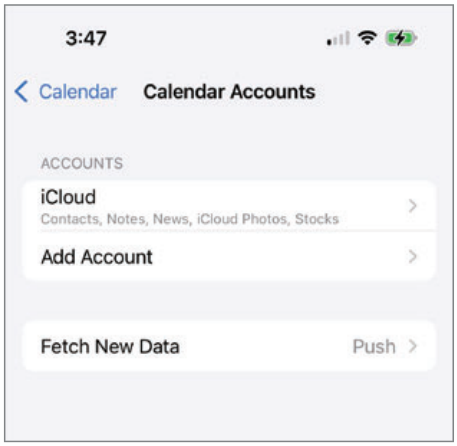

Android devices will, by default, sync calendar and contacts with the user's `@gmail.com` ID. For example, if you go to Settings ➢ Apps ➢ Calendar ➢ Calendar Settings ➢ Manage Calendars, you will get a screen like the one in Figure 10.65. Here the toggle is on to sync with the Google account. To add a different account to sync to, tap the + sign in the upper-right corner and add the account.

FIGURE 10.65 Android Manage Calendars

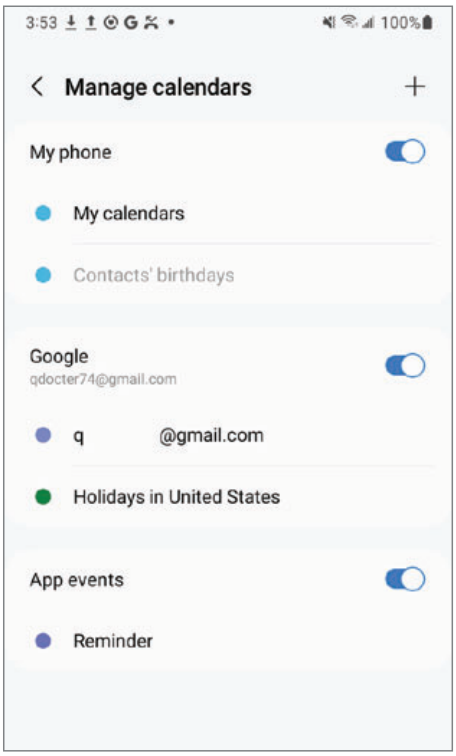

Each mobile OS has an option to manually force synchronization if the user wants to. iOS generally uses the language "Fetch New Data," whereas Android prefers "Sync Now." Invoking either of those options will manually force a sync for whatever app the user is in.

Microsoft 365

Microsoft 365 is a subscription service that provides access to the Office suite of apps from Microsoft, including Word, Excel, PowerPoint, and others. With the subscription, users also get storage space in Microsoft's cloud. That cloud storage space can also be used to sync devices with one another.

Imagine a scenario where a user has a desktop and laptop computer, and they want to ensure the operating environment is always identical between the two. Windows users with a Microsoft 365 account can easily sync settings between multiple computers with their

subscription. Simply open Start ➢ Settings ➢ Accounts ➢ Windows Backup (Figure 10.66). Here there are options to sync folders and files on OneDrive, apps, as well as preferences such as accessibility, accounts, and personalization features such as the desktop. For this to work, synchronization needs to be enabled on each computer the user wants to synchronize. Remember that the synchronization will be based on the user's email address registered with Microsoft.

FIGURE 10.66 Windows Backup Sync Settings

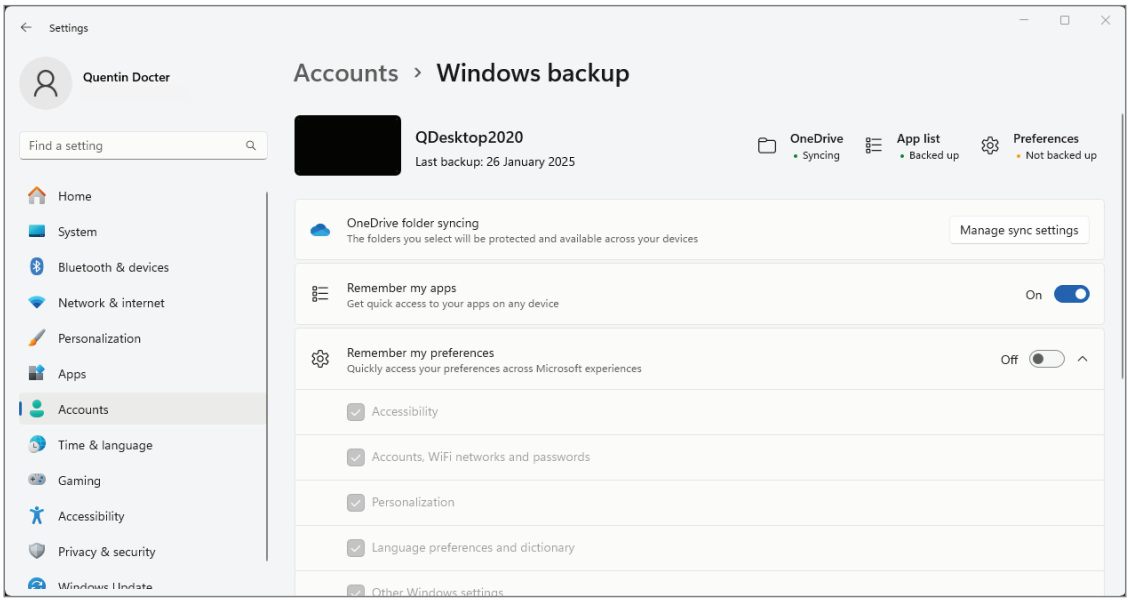

If synchronization is no longer desired, the user can remove the settings from the cloud by taking the following steps:

1. Turn off Sync settings on all synced computers.
2. Go to `https://account.microsoft.com/devices` and log in.
3. Scroll to the bottom of the page and select Clear Stored Settings.

To synchronize mobile devices with OneDrive, download the OneDrive app from the appropriate app store.

Using Cloud Storage

Storage is the area in which cloud computing got its start. The idea is simple—users store files just as they would on a hard drive, but on servers accessed over the Internet. It makes it so different users can access the files regardless of where they are physically located. Users can be located in the United States, China, and Germany; they all have access via their web browser. This is particularly helpful for multinational organizations.

Exam Objective 1.3 includes synchronizing through cloud storage. Understand what cloud storage is and be able to provide a few examples.

There is no shortage of cloud-based storage providers in the market today. Each provider offers slightly different features. Most of them will offer limited storage for free and premium services for more data-heavy users. Table 10.6 shows a comparison of some of the better-known providers. Please note that the data limits and prices can change; this table is provided for illustrative purposes only and doesn't include every level of premium service available. Most of these providers offer business plans with unlimited storage as well for an additional cost.

TABLE 10.6 Cloud providers and features

Service	Free	Premium	Cost per year
Google Drive	15 GB	2TB	$120
Apple iCloud	5 GB	2TB	$120
Microsoft OneDrive	5 GB	1TB	$100
Dropbox	2 GB	3TB	$199
Box	10 GB	100 GB	$60

Which one should you choose? If you want extra features, such as web-based applications or artificial intelligence (AI), then Google or Microsoft are probably the best choices. If you just need data storage, then another option might be better. Some allow multiple users to access a personal account, so that might figure into your decision as well.

Most cloud storage providers offer synchronization to the desktop, which makes it so that you have a folder on your computer, just as if it were on your hard drive. And it's important to note that folder will almost always have the most current edition of the files stored in the cloud. The *synchronization app* typically runs in the background and has configurable options, including what, when, and how often to synchronize.

Accessing the sites is done through your web browser. Once you are in the site, managing your files is much like managing them on your local computer. In Figure 10.67, you can see the Google Drive interface, with a few files and folders in it.

Sharing files with others is easy. Right-click the file or folder and choose Share. You'll be asked to enter their name or email address and to indicate whether anyone on the Internet with the link can view it, or only people you specify. For security reasons, it's better to only allow people you specify to view the file or folder.

FIGURE 10.67 Google Drive

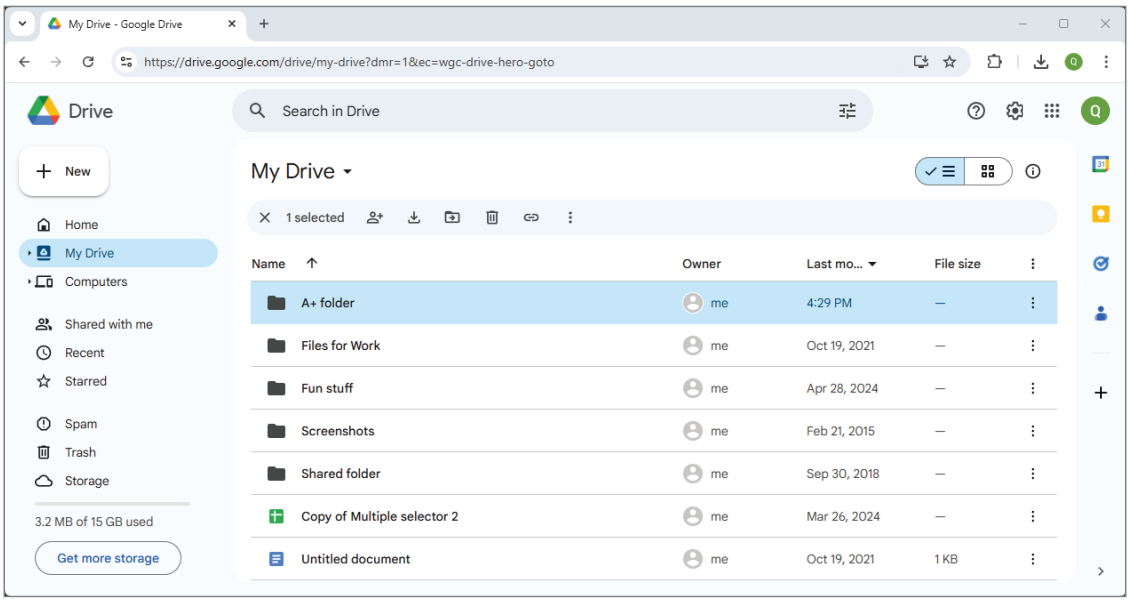

Summary

This chapter introduced you to key features of mobile devices: network connectivity, device management, and synchronization. Establishing network connectivity means enabling cellular, Wi-Fi, or Bluetooth connections and configuring them properly.

Key cellular concepts to understand are wireless standards, hotspots and tethering, and SIM/eSIM. When using Bluetooth, you need to pair the device with another Bluetooth device for connectivity.

Next, we covered mobile app support. A popular feature of mobile devices is location services, which can help locate you or your phone and give directions. We then moved into mobile device management and application management, including corporate and BYOD device configurations, policy enforcement, and corporate applications with particular emphasis on email configuration.

Finally, we looked at synchronization. Synchronizing a mobile device to the cloud or a desktop/laptop is a good way to ensure that data is saved to a secure or permanent location. Types of data that are synchronized include calendars and contacts, as well as business applications such as email and cloud storage. Also understand how to recognize and manage data caps. Common connection types for synchronization include Wi-Fi, USB, and cellular.

Exam Essentials

Understand the differences between wireless specifications. Know basic differences between 3G, 4G, and 5G, including what GSM and CDMA are and why they weren't compatible.

Know what a hotspot is. A hotspot is a cellular-capable device such as a smartphone, that is sharing its cellular connection with other devices so they can access the Internet.

Know how to enable or disable wireless/cellular data connections. Wireless connections are individually disabled through the Settings app on most phones or through a quick access screen. Airplane mode disables all wireless signals at once, when configured properly.

Know what SIM and eSIM are. The Subscriber Identity Module (SIM) is the chip in the mobile device that identifies it. An eSIM is a circuit embedded into the device's motherboard that does the same function, but can be electronically reprogrammed, much like a BIOS/UEFI can be flashed.

Understand the steps needed to configure Bluetooth. You need to enable Bluetooth, enable pairing, find a device for pairing, enter the appropriate PIN code (or confirm the PIN), and test connectivity.

Know the differences between GPS and cellular location services. GPS is a free service provided by the government and uses satellites. Cellular location services use cell phone towers and require an account with a carrier.

Know the purposes of MDM and MAM. Mobile device management (MDM) is primarily used to determine which mobile devices are allowed on a network and to set policies for access. MDM also provides mechanisms for remotely locking and wiping devices. Mobile application management (MAM) is for managing corporate applications on mobile devices.

Know which protocols are used for email and which ports they use. POP3 (port 110) and IMAP (port 143) are used to receive email, and SMTP (port 25) is used to send email. None of these protocols are inherently secure. They can be secured with SSL or TLS. SMTP over SSL uses port 465, and SMTP over TLS uses port 587. IMAP over SSL/TLS uses port 993, and POP over SSL/TLS uses port 995.

Understand the differences between managing corporate devices and BYOD devices. Corporate devices are owned by the company or organization, and all content on them is considered property of that company. BYOD devices are owned by the end user and allowed to participate on the company network. BYOD configurations can be less secure because there can be multiple device types and OSs, but using a BYOD policy can save companies money.

Be familiar with four commercial email providers and required configuration items. Common commercial email providers are Gmail, iCloud, Exchange Online, and Yahoo Mail. Each provider has its own inbound and outbound servers, but most of the time that configuration information is automatically provided when you try to connect to them with an email client.

Know which types of data are often synchronized. Common data types for synchronization include calendars and contacts, but literally any data on a mobile device can be synchronized. This includes business applications such as email, and cloud storage.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. A friend asks you why she would ever disable cellular data networking. Which of the following is *not* a valid reason to disable cellular data networking?
   - **A.** You have a limited amount of data in your monthly plan.
   - **B.** You have access to a reliable Wi-Fi signal.
   - **C.** You are about to download an update to your phone.
   - **D.** Your phone calls are going out over your carrier's cellular network.
2. You need to pair two Bluetooth devices with each other. Which of the following is true regarding Bluetooth pairing?
   - **A.** Bluetooth connections support wireless device control but not file transfers.
   - **B.** Bluetooth is not yet a fully standardized protocol.
   - **C.** Bluetooth connections do not reach as far as Wi-Fi connections.
   - **D.** You must reboot the paired devices to complete the connection.
3. A technician is configuring a mobile client to use email through an ISP. Which port number should he configure to use SMTP?
   - **A.** 25
   - **B.** 110
   - **C.** 143
   - **D.** 995
4. You need to configure email for a mobile client. His device has a setting referring to port 587. Which mail protocol is likely using this protocol?
   - **A.** SMTP with SSL
   - **B.** SMTP with TLS
   - **C.** IMAP4 with SSL/TLS
   - **D.** POP3 with SSL/TLS
5. You are enabling email on a mobile client. Which of the following is *not* a standard integrated commercial provider for email?
   - **A.** iCloud
   - **B.** Hotmail
   - **C.** Yahoo Mail
   - **D.** Exchange Online

6. If a user wanted to disable location services on an iPhone, where would they do that?
   - **A.** Settings ➢ Location Services
   - **B.** Settings ➢ Communications
   - **C.** Settings ➢ Privacy & Security
   - **D.** Settings ➢ GPS

7. A mobile device is having intermittent communications issues. Which software component of a mobile phone is most likely responsible for these issues?
   - **A.** Either iOS or Android OS
   - **B.** SIM OS
   - **C.** Baseband OS
   - **D.** Wireless OS

8. An administrator needs to be able to control which apps are installed on the corporate network's mobile devices, and update or delete them and their data as necessary. Which type of service does she need?
   - **A.** MAM
   - **B.** MDM
   - **C.** Location services
   - **D.** Device synchronization

9. You need to synchronize Windows desktop settings between two laptop computers. Which service is designed to do this?
   - **A.** Microsoft 365
   - **B.** MAM
   - **C.** MDM
   - **D.** iCloud

10. You are pairing a Bluetooth headset to a mobile device. On the mobile device, after finding the device for pairing, what is the next step to take?
    - **A.** Enable Bluetooth.
    - **B.** Test connectivity.
    - **C.** Enter the appropriate PIN code.
    - **D.** Enable pairing.

11. Which of the following most accurately describes the purpose of a SIM?
    - **A.** It's the data storage space for a mobile device.
    - **B.** It stores phone identification and security data.
    - **C.** It manages the mobile device's wireless connections.
    - **D.** It's the controller for location services on a mobile device.

12. Which of the following location services typically provides the best accuracy when detecting a user's location?
    - **A.** Cellular
    - **B.** GPS
    - **C.** eSIM
    - **D.** MDM

13. Which of the following statements about configuring email access on a mobile device is true?
    - **A.** Most Internet mail services offer an Exchange option.
    - **B.** The TCP ports used for configuring access are usually standard port numbers.
    - **C.** Most ports used for access are UDP ports.
    - **D.** You must download third-party apps for connecting to email services.

14. Which of the following are disabled on an Android device when airplane mode is switched on? (Choose all that apply.)
    - **A.** Cellular
    - **B.** Wi-Fi
    - **C.** Bluetooth
    - **D.** Settings app
    - **E.** Lock screen

15. You need to enable clients to synchronize their mobile devices to their laptops. Which of the following are common connection types to enable laptop synchronization? (Choose all that apply.)
    - **A.** Cellular
    - **B.** Bluetooth
    - **C.** Wi-Fi
    - **D.** USB
    - **E.** Tethering

16. A user calls in from their morning commute, using their iPad, which seems to be experiencing data problems. While traveling in the car or on the train, they can't download data from the corporate cloud. However, they can once they are in the office. This has been happening for three days. What is the most likely cause of the problem?
    - **A.** The train is moving too fast for the cellular connection to be able to download data.
    - **B.** The cellular connection is disabled. Check to ensure that airplane mode is not enabled.
    - **C.** The user has reached the cellular data cap for the month.
    - **D.** Their iPad is configured as BYOD and needs to be switched to corporate.

17. When changing an Android smartphone from one service provider to another, which chip would be physically removed from the device and replaced with a new one?

A. SIM

B. eSIM

C. GPS

D. BYOD

18. A network administrator wants to be able to remotely wipe everything off of corporate smartphones if they are lost or stolen. What type of service does she need?

A. MAM

B. MDM

C. eSIM

D. Location services

19. A user needs to synchronize her iPhone to another platform. Which two options are the default ones for her to choose? (Choose two.)

A. Google Drive

B. SharePoint

C. iCloud

D. iTunes

20. You are investigating potential location services options for a client. Which type of location services will need to be paid for versus being free of charge?

A. GPS

B. Cellular

C. Bluetooth

D. Wi-Fi

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answers compare to the authors', refer to Appendix B.

Explain how to establish Wi-Fi connectivity on an Apple iPhone.

Chapter 11

Virtualization and Cloud Computing

THE FOLLOWING COMPTIA A+ 220-1201 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

✓ **4.0 Virtualization and Cloud Computing**

✓ **4.1 Explain virtualization concepts.**

- Purpose of virtual machines
  - Sandbox
  - Test development
  - Application virtualization
    - Legacy software/OS
    - Cross-platform virtualization
- Requirements
  - Security
  - Network
  - Storage
- Desktop virtualization
  - Virtual Desktop Infrastructure (VDI)
- Containers
- Hypervisors
  - Type 1
  - Type 2

✓ **4.2 Summarize cloud computing concepts.**

- Common cloud models
  - Private cloud
  - Public cloud

- Hybrid cloud
- Community cloud
- Infrastructure as a service (IaaS)
- Software as a service (SaaS)
- Platform as a service (PaaS)

- Cloud characteristics
  - Shared resources vs. dedicated resources
  - Metered utilization
    - Ingress/egress
  - Elasticity
  - Availability
  - File synchronization
  - Multitenancy

The computer industry is one of big trends. A new technology comes along and becomes popular, until the next wave of newer, faster, and shinier objects comes along to distract everyone from the previous wave. Thinking back over the past 20 to 30 years, there have been several big waves, including the rise of the Internet, wireless networking, and mobile computing.

Within each trend, there are often smaller ones. For example, the Internet was helped by modems and ISPs in the middle of the 1990s, and then broadband access took over. Wireless networking has seen several generations of faster technology, from the 11 Mbps 802.11b, which at the time it came out was pretty cool, to 802.11ax and newer, delivering gigabit and faster wireless. Mobile computing has been a long-lasting wave, first with laptop computers becoming more popular than desktops, and then with handheld devices (namely, smartphones and tablets) essentially functioning like computers.

The biggest recent wave in the computing world is cloud computing. Its name comes from the fact that the technology is Internet-based; in most computer literature, the Internet is represented by a graphic that looks like a cloud.

Traditionally, the delivery method for services has been that the servers are on the same network as the clients. They might not be on the same LAN, but they are certainly administered by one company or one set of administrators. If clients on the network need a new feature, the network architects and administrators add the necessary server. This is still the most common setup today, but cloud computing makes that model almost obsolete. In essence, cloud computing lets networks break out of that model and have services provided by a server that the company doesn't own, and so it's not under the company's direct control. Behind the scenes, virtualization is the most important technology in cloud computing because it removes the barrier of needing one-to-one relationships between the physical computer and the operating system.

There's been a sharp growth in cloud computing and virtualization in the last several years. It seems like everyone is jumping on the cloud (pun intended, but doesn't that sound like fun?), and technicians need to be aware of what it can provide and its limitations. We will cover both topics in this chapter, starting with the underlying virtualization technology and then moving into a much grander and fluffier-sounding application.

Understanding Virtualization

Perhaps the easiest way to understand virtualization is to compare it to more traditional technologies. In the traditional computing model, a computer is identified as being a physical

machine that is running some combination of software, such as an operating system and various applications. There's a one-to-one relationship between the hardware and the operating system.

For the sake of illustration, imagine that a machine is a file server and now it needs to perform the functions of a web server as well. To make this happen, the administrator would need to ensure that the computer has enough resources to support the service (CPU, memory, network bandwidth), install web server software (Apache HTTP Server or Microsoft Internet Information Services [IIS], for example), configure the appropriate files and permissions, and then bring it back online as a file and web server. These would be relatively straightforward administrative tasks.

But now imagine that the machine in question is being asked to run Windows Server and Linux at the same time. Now there's a problem. In the traditional computing model, only one OS can run at one time, because each OS completely controls the hardware resources in the computer. Sure, an administrator can install a second OS and configure the server to dual-boot, meaning the OS to run is chosen during the boot process, but only one OS can run at a time. So if the requirement is to have a Windows-based file server and a Linux-based Apache web server, there's a problem. Two physical computers are needed.

Similarly, imagine that there is a Windows-based workstation being used by an applications programmer. The programmer has been asked to code an app that works in Linux, or Apple's iOS, or anything other than Windows. When the programmer needs to test the app to see how well it works, what do they do? Sure, they can configure their system to dual-boot, but once again, in the traditional computing model, they are limited to one OS at a time per physical computer. Their company could purchase a second system, but that quickly starts to get expensive when you have multiple users with similar needs.

This is where virtualization comes in. The term *virtualization* is defined as creating virtual (rather than actual) versions of something. In computer jargon, it means creating virtual environments where "computers" can operate. We use quotation marks around the word *computers* because they don't need to be physical computers in the traditional sense. Virtualization is often used to let multiple OSs (or multiple instances of the same OS) run on one physical machine at the same time. Yes, they are often still bound by the physical characteristics of the machine on which they reside, but virtualization breaks down the traditional one-to-one relationship between a physical set of hardware and an OS.

Virtualization has been around in the computer industry since 1967, but it has only recently exploded in popularity thanks to the flexibility that the Internet offers.

The Purpose of Virtual Machines

We have already hit on the major feature of virtualization, which is breaking down that one-to-one hardware and software barrier. The virtualized version of a computer is appropriately called a *virtual machine (VM)*. Thanks to VMs, it is becoming far less common to need dual-boot machines today than in the past. In addition, VMs make technology like the

cloud possible. A cloud provider can have one incredibly powerful server that is running five instances of an OS for client use, and each client is able to act as if it had its own individual server. On the flip side, cloud providers can pool resources from multiple physical servers into what appears to the client to be one system, effectively giving clients unlimited processing or storage capabilities (assuming, of course, that the provider doesn't physically run out of hardware).

Virtual OSs can be powered on or off individually without affecting the host OS or hypervisor—we'll talk about what those are in the upcoming section "The Hypervisor."

Virtual machines have a wide variety of applications, many of which are cloud-based services. Here are three specific uses you should pay particular attention to:

Virtual Sandbox Imagine a scenario where you have an application that you want to test out in an OS, but you don't want any negative effects to happen to the computer system doing the testing. One way to do this is to test the app in a *sandbox*, which is a temporary, isolated desktop environment. Think of it as a temporary, somewhat limited virtual machine. Any app operating in the sandbox will act as it would in a full version of the chosen OS, with one big difference. Files are not saved to the hard drive or memory, so the physical machine should never be affected by anything the app in the sandbox does. When the sandbox gets shut down, the app and any data associated with it goes away.

Exam Objective 4.1, "Explain virtualization concepts," wants you to be able to explain the purpose of virtual machines. This includes sandbox, test development, and application virtualization (for legacy software or OSs and cross-platform virtualization).

There are several sandboxing software solutions on the market, including Sandboxie-Plus, SHADE Sandbox, VirtualBox, CodeSandbox, and ToolWiz TimeFreeze. Some of them are designed for app development or cloud testing, whereas others will literally sandbox your whole system unless you authorize changes to specific files on the computer. Last but not least, Microsoft is in on the action with Windows Sandbox as well.

Windows Sandbox can be installed through the Turn Windows Features On Or Off app. Type **`windows features`** into the Windows 11 search bar, and the app name should appear on the right side of the search window. Open it, and in the list of features scroll down to Windows Sandbox. Select the option, click OK, and reboot if prompted. If the Windows Sandbox option is not available, then the computer doesn't meet the requirements needed to run Sandbox.

Test Development We already gave an example of *test development* in a virtual machine, with a developer looking to create and test applications as they are works in progress. Apps can be tested in virtualized versions of the same OS that the developer is working on, or in a different OS virtualized on the same computer.

Application Virtualization *Application virtualization* is a common use of virtual machines as well. It usually takes one of two forms. The first is virtualizing *legacy software* or a legacy OS. Legacy software is basically old, outdated software. The problem is that legacy apps often only run on legacy OSs, so you need to either virtualize the app in a newer OS and tweak it like crazy to get it to work or virtualize it in an older OS that has no business running on its own server either.

The second use is *cross-platform virtualization*. It allows programs coded for one type of hardware or operating system to work on another that it's not designed to work on. For example, an app designed for macOS could work in a virtualized version of that OS within a Windows-based server.

The underlying purpose for these types of virtual machines is to save money. Administrators don't need to deal with older hardware or software, which inherently pose bigger security risks than newer versions. And if an older, business-critical app is still working well enough within a virtualized environment, then it falls into the "if it isn't broke, don't fix it" category. End users, in the workstation example we provided earlier, can have multiple environments to use without needing to buy additional hardware as well.

The Hypervisor

The key enabler for virtualization is a piece of software called the *hypervisor*, also known as a *virtual machine manager (VMM)*. The hypervisor software allows multiple operating systems to share the same host, and it also manages the physical resource allocation to those virtual OSs. As illustrated in Figure 11.1, there are two types of hypervisors: Type 1 and Type 2.

Exam Objective 4.1 does expect you to know the difference between Type 1 and Type 2 hypervisors.

A *Type 1 hypervisor* sits directly on the hardware, and because of this, it's sometimes referred to as a *bare-metal hypervisor*. In this instance, the hypervisor is basically the operating system for the physical machine. This setup is most commonly used for server-side virtualization, because the hypervisor itself typically has very low hardware requirements to support its own functions. Type 1 is generally considered to have better performance than Type 2, simply because there is no host OS involved and the system is dedicated to supporting virtualization. Virtual OSs are run within the hypervisor, and the virtual (guest) OSs are completely independent of each other. Examples of Type 1 hypervisors include Microsoft

Hyper-V, VMware ESXi, and Citrix Hypervisor (formerly XenServer). Exercise 11.1 walks you through the steps to enable Hyper-V in Windows 11.

FIGURE 11.1 Type 1 and Type 2 hypervisors

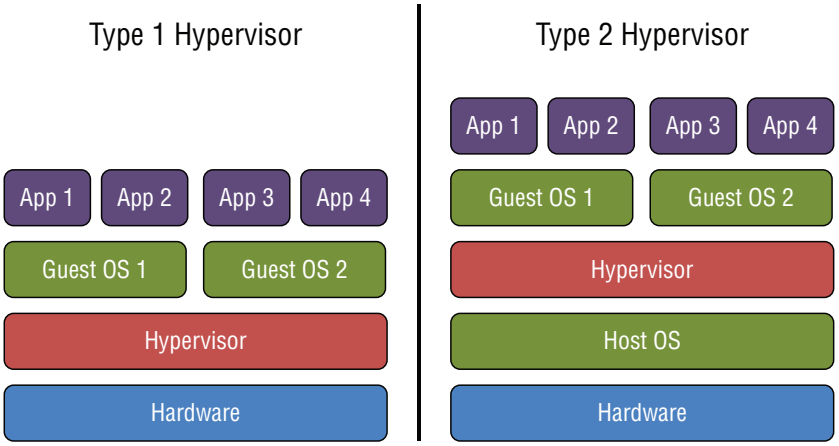

EXERCISE 11.1

Enabling Hyper-V in Windows 11

To enable Microsoft Hyper-V, perform the following steps:

1. Check for minimum system requirements.
   - Windows 11 Enterprise or Pro
   - 64-bit processor with Second Level Address Translation (SLAT)
   - CPU support for VM Monitor Mode Extension (VT-c on Intel CPUs)
   - 4 GB of RAM or more
   - Hardware virtualization support enabled in the BIOS/UEFI
2. If requirements are met, then open the Turn Windows Features On Or Off app. Start typing the name of the app into the Windows search box and press Enter when it appears on the right side. The Windows Features app is shown in Figure 11.2.

EXERCISE 11.1 *(continued)*

FIGURE 11.2 Enabling Hyper-V

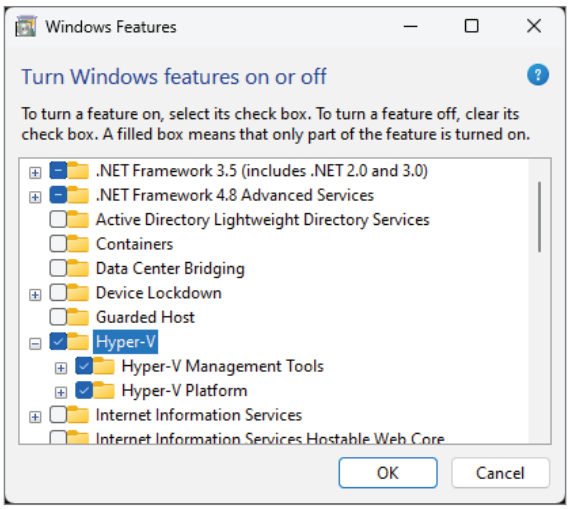
Windows Features
Turn Windows features on or off
To turn a feature on, select its check box. To turn a feature off, clear its check box. A filled box means that only part of the feature is turned on.
.NET Framework 3.5 (includes .NET 2.0 and 3.0)
.NET Framework 4.8 Advanced Services
Active Directory Lightweight Directory Services
Containers
Data Center Bridging
Device Lockdown
Guarded Host
Hyper-V
Hyper-V Management Tools
Hyper-V Platform
Internet Information Services
Internet Information Services Hostable Web Core
OK
Cancel

3. Select the Hyper-V option and click OK.
4. After installation, restart your computer.
5. After restarting, you will be able to open the Hyper-V Manager (Figure 11.3) by typing **`Hyper-V`** into the Windows search box and pressing Enter.

FIGURE 11.3 Hyper-V Manager

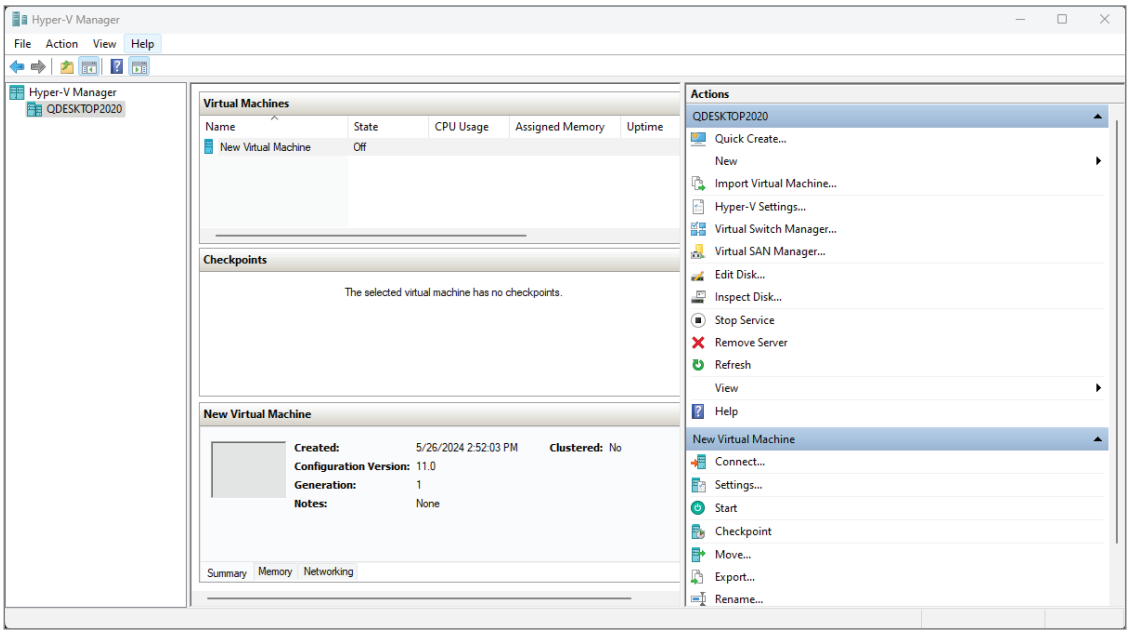

A *Type 2 hypervisor* sits on top of an existing operating system, called the *host OS*. This is most commonly used in *client-side virtualization*, where multiple OSs are managed on the client machine as opposed to on a server. An example of this would be a Windows user who wants to run Linux at the same time as Windows. The user could install a hypervisor, then install Linux as a guest OS in the hypervisor and run both OSs concurrently and independently. The downsides of Type 2 hypervisors are that the host OS consumes resources, such as processor time and memory, and a host OS failure means that the guest OSs fail as well. Examples of Type 2 hypervisors include Microsoft's Windows Virtual PC and Azure Virtual Server, Oracle VM VirtualBox, VMware Workstation, and Linux KVM.

Containers

Hypervisors and the virtualization service they provide are wonderful, but sometimes a smaller and faster solution is needed for just one app. That solution is a container. A *container* is an application and all of the OS files needed to run it (called runtime dependencies) virtualized into one package. Think of it as virtualizing the app and the OS together to ensure a consistent app experience across different OSs that could potentially run the app. Unlike a hypervisor, it just virtualizes the software as opposed to the underlying physical infrastructure. An example of what this logically looks like compared to a virtual machine is shown in Figure 11.4.

FIGURE 11.4 Container versus virtual machine

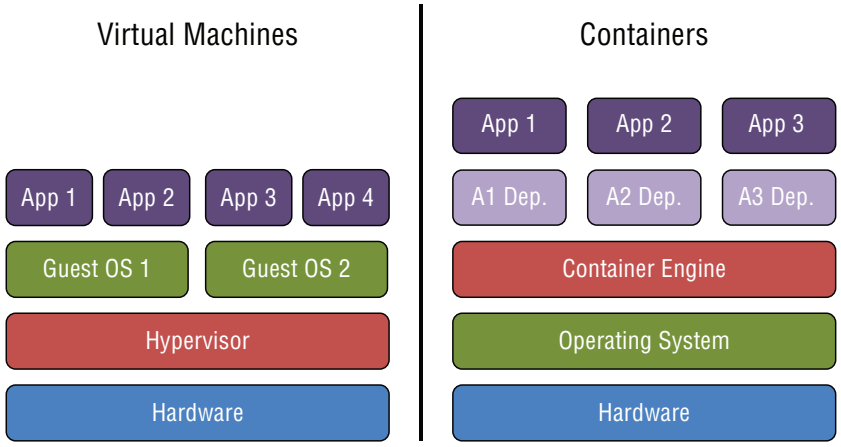

To use containers, you need a containerization software package, or a *container engine*. The original and still popular Docker is one example, but there are others such as Buildah, Podman, LXD, and Containerd. The software developer would create the *container image*, a virtualization ISO file that includes the app and all dependencies and environment files needed to run the app. Then, that file can be distributed to other users for use on their computers, regardless of the operating system. On the receiving end, a container engine running as an app in the OS executes and manages the container. Docker, Kubernetes, LXD, and others are used here.

Whereas cross-platform app development may be better suited for a virtual machine, once an app is developed, deployment of it across a large network or even the cloud can be a lot faster using containers. Containers are also smaller in size than virtual machines and much easier to scale to large or small network environments. Table 11.1 summarizes the differences between containers and virtual machines.

TABLE 11.1 Comparing containers and virtual machines

	Container	Virtual Machine
What is it?	Package of software with an application and its dependencies	A digital replica of a physical computer
What does it virtualize?	An app in an OS runtime environment	Physical infrastructure
How many applications?	Usually one, but can be more than one	Unlimited (only limited by hardware resources)
Package size	Usually megabytes	Usually gigabytes
Scope of control	Container only	All hardware and software in the VM
Speed of deployment	Fast	Slower
Scalability	Highly scalable	Can be slow and expensive to scale

Client-Side Virtualization Requirements

As you might expect, running multiple OSs on one physical workstation can require more resources than running a single OS. There's no rule that says a workstation being used for virtualization is required to have more robust hardware than another machine, but for performance reasons, the system should be well equipped. This is especially true for systems running a Type 2 hypervisor, which sits on top of a host OS. The host OS will need resources, too, and it will compete with the VMs for those resources. Let's talk about specific requirements.

Resource Requirements

The primary resources here are the same as you would expect when discussing any other computer's performance: CPU, RAM, hard drive storage space, and network performance. From the CPU standpoint, know that the hypervisor can treat each core of a processor as a separate virtual processor, and it can even create multiple virtual processors out of a single core. The general rule here is that the faster the processor the better, but really, the more

cores a processor has, the more virtual OSs it can support in a speedy fashion. Within the hypervisor, there will most likely be an option to set the allocation of physical resources, such as CPU priority and amount of RAM, to each VM.

Some hypervisors require that the CPU be specifically designed to support virtualization. For Intel chips, this technology is called *virtualization technology (VT)*, and AMD chips need to support *AMD-V*. Pretty much every processor today supports virtualization, but you might run across older ones that do not. In addition, many system BIOSs/UEFIs have an option to turn on or turn off virtualization support. If a processor supports virtualization but the hypervisor won't install, check the BIOS/UEFI and enable virtualization. The specific steps to do this vary based on the BIOS/UEFI, so check the manufacturer's documentation. An example of what this might look like is shown in Figure 11.5.

FIGURE 11.5 Enabling virtualization in the BIOS/UEFI

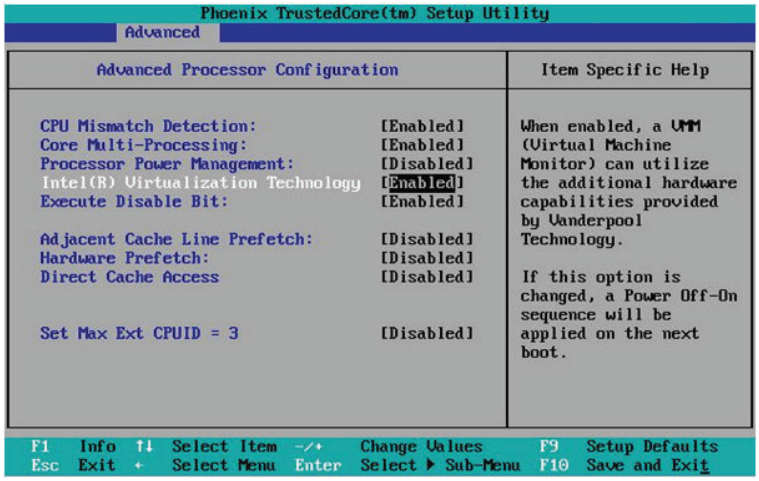
Phoenix TrustedCore(tm) Setup Utility
Advanced
Advanced Processor Configuration
Item Specific Help
CPU Mismatch Detection: [Enabled]
Core Multi-Processing: [Enabled]
Processor Power Management: [Disabled]
Intel(R) Virtualization Technology [Enabled]
Execute Disable Bit: [Enabled]
Adjacent Cache Line Prefetch: [Disabled]
Hardware Prefetch: [Disabled]
Direct Cache Access [Disabled]
Set Max Ext CPUID = 3 [Disabled]
When enabled, a VMM (Virtual Machine Monitor) can utilize the additional hardware capabilities provided by Vanderpool Technology.
If this option is changed, a Power Off-On sequence will be applied on the next boot.
F1 Info ↑↓ Select Item -/+ Change Values F9 Setup Defaults
Esc Exit ← Select Menu Enter Select ▶ Sub-Menu F10 Save and Exit

Memory is always a big concern for computers, and virtual ones are no different. When you're installing the guest OS, the hypervisor will ask how much memory to allocate to the VM. This can be modified later if the guest OS in the VM requires more memory to run properly. Always remember, though, that the host OS requires RAM, too. Thus, if the host OS needs 4 GB of RAM and the guest OS needs 4 GB of RAM, the system needs to have at least 8 GB of RAM to support both adequately.

Hard disk space works the same way as RAM. Each OS will need its own hard disk space, and the guest OS will be configured via the hypervisor. Make sure that the physical computer has enough free disk space to support the guest OSs.

Exam Objective 4.1 includes network and storage requirements for virtualization. Be sure to understand that each virtual OS needs its own resources on the physical machine on which they reside.

Finally, from a networking standpoint, each of the virtual desktops will typically need full network access, and configuring the permissions for each can sometimes be tricky.

The virtual desktop is often called a *virtual desktop infrastructure (VDI)*, a term that encompasses the software and hardware needed to create the virtual environment. VDIs can be on premises in the same building as the company using it, or in the cloud. The VM will create a *virtual NIC* and manage the resources of that NIC appropriately. The virtual NIC doesn't have to be connected to the physical NIC (unless access to other computers across a physical network, such as the Internet, is required); an administrator could create an entire virtual network within the virtual environment where the virtual machines just talk to each other.

Desktop virtualization, and the term VDI, are included in Exam Objective 4.1. Be able to explain these terms!

That's not normally practical in the real world, though, so the virtual NIC will be connected to the physical NIC. Configuring a virtual switch within the hypervisor normally does this. The virtual switch manages the traffic to and from the virtual NICs and logically attaches to the physical NIC (see Figure 11.6). Network bandwidth is often the biggest bottleneck when you are running multiple virtual OSs. If the network requirements are that each of four VMs on one physical machine get Gigabit Ethernet service all the time, one physical Gigabit Ethernet NIC isn't going to cut it.

FIGURE 11.6 Virtual NICs connecting to a physical NIC

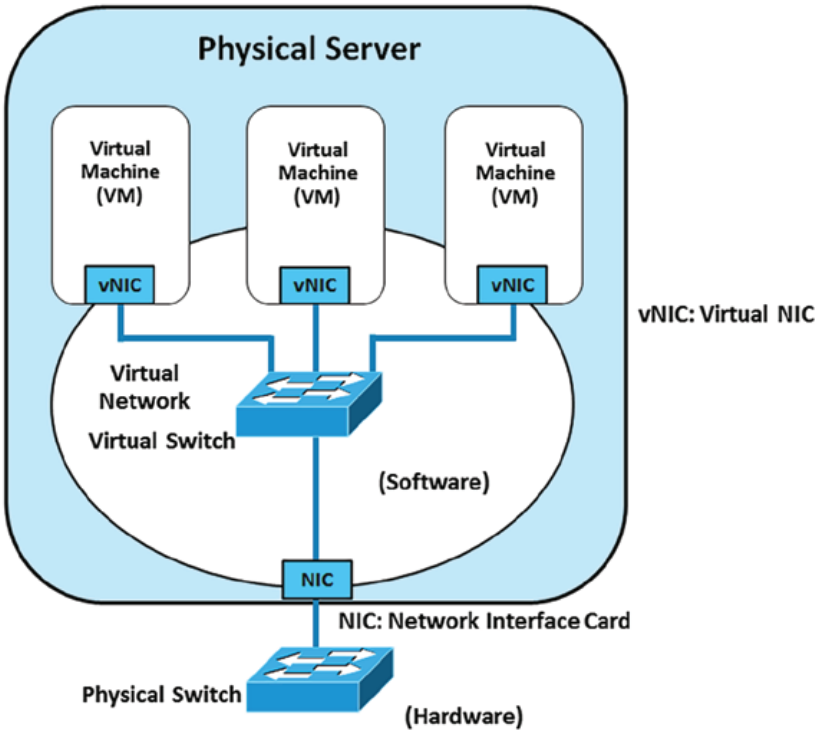

With Type 1 hypervisors, virtual desktops are often used with remote administration. This can allow a remote administrator to work on the workstation (to perform maintenance, for example) with or without the knowledge of the user sitting in front of the machine.

Emulator Requirements

Virtual machines are created to exist and function just like a physical machine. Thus, all the requirements that a physical machine would have need to be replicated by the hypervisor, and that process is called *emulation*. The terms *hypervisor* and *emulator* are often used interchangeably, although they don't mean the same thing. The hypervisor can support multiple OSs, whereas technically, an emulator appears to work the same as one specific OS. As for requirements, the emulator and the hypervisor need to be compatible with the host OS. That's about it.

Some incredibly popular mobile-based games do not have PC equivalents. In other words, you need to run them from iOS or Android. Android is based on open source code, meaning that there are Android emulators on the market. One such free emulator is Andy (`www.andyroid.net`). People can install Andy on their desktops or laptops, and then install Android apps within Andy. That way, they can play their favorite mobile games on their desktop or laptop computers.

Security Requirements

In the early days of the cloud, a common misconception was that virtual machines couldn't be hacked. Unfortunately, some hackers proved this wrong. Instead of attacking the OS in the VM, hackers have turned their attention to attacking the hypervisor itself. Why just hit one OS when you can hit all of them on the computer at the same time? A number of virtualization-specific threats focusing on the hypervisor have cropped up, but updates have fixed the issues as they have become known. The solution to most virtual machine threats is to always apply the most recent updates to keep the system(s) current.

Virtualization security requirements are included in Exam Objective 4.1. Be sure how to understand how security works in an emulated environment.

At the same time, all the security concerns that affect individual computers also apply to VMs. For example, if Windows is being operated in a VM, that instance of Windows still needs antimalware software installed on it. As a rule of thumb in a cloud, the vendor is responsible for securing the hardware and network connections, whereas the client is responsible for securing the data stored in the cloud.

Real World Scenario

Setting Up and Using Client-Side Virtualization

If given a scenario on the A+ exam or faced with the situation in real life, you should be able to set up and show someone how to use client-side virtualization. Here are some sample steps to follow:

1. Determine the client's needs for virtualization.

   For example, do they have a Windows or Mac client and need to run Linux? Or perhaps they have a Windows computer and want to run macOS at the same time? Determine their needs, and then secure the additional OSs (including licenses, as appropriate) before beginning the installation.

2. Evaluate the computer to ensure that it can support the VM.
   - Does the processor support virtualization?
   - How much RAM does the computer have? Is it enough to meet the minimum requirements of all installed OSs?
   - How much free hard drive space is there? It should be enough to install the hypervisor and the guest OS as well as to store any files that need to be stored from within the guest OS.
   - Does the system have a fast enough network connection if the guest OS needs access to the network?
3. Consider which hypervisor to use. Is there a version of the hypervisor that's compatible with the host OS?
4. Consider security requirements. If the guest OS will be on the Internet, will proper security software be installed?
5. After all conditions are deemed sufficient, you can install the hypervisor and the guest OS.

It will not affect the host OS, but it's always a good idea to back up the system before installing any new major software packages!

Exercise 11.2 will give you hands-on experience installing a hypervisor and a guest OS.

Now that we have covered the key concepts behind client-side virtualization, it's time to practice. Exercise 11.2 walks you through installing the Oracle VirtualBox hypervisor on a Windows 11 computer and then installing Lubuntu (a distribution of Linux). Normally, installing a second OS involves a relatively complicated process where you need to dual-boot

your computer. You're not going to do that here. Instead, you will use the VirtualBox hypervisor that allows you to create a new virtual system on your hard drive and not affect your existing Windows installation. We promise you that this exercise will not mess up Windows on your computer! And when you're finished, you can just uninstall VirtualBox, if you want, and nothing will have changed on your system. This exercise is admittedly a bit long because there are a lot of steps, and it's also probably more "advanced" than typical A+ materials. That said, we encourage you to try it—it usually ends up being one of our students' favorite exercises during training classes.

EXERCISE 11.2

Installing VirtualBox and Lubuntu on Windows 11

The first two steps are for preparation. You need to download Oracle VirtualBox and a version of Lubuntu. Really, any version of Linux is fine, but we'll point you to Lubuntu, which is a good choice. Depending on your network speed, the download could take a bit. These directions are based on VirtualBox version 7.1.6. If you have a different version, the wording on some of the screens might be slightly different, but the following steps should still work.

Obtaining VirtualBox and Lubuntu

1. Download Oracle VirtualBox from `www.virtualbox.org/wiki/Downloads` (see Figure 11.7). Select VirtualBox For Windows Hosts, unless, of course, you have a Mac, and then you need the one for macOS hosts. Save it to your desktop for ease of access.

FIGURE 11.7 Download VirtualBox

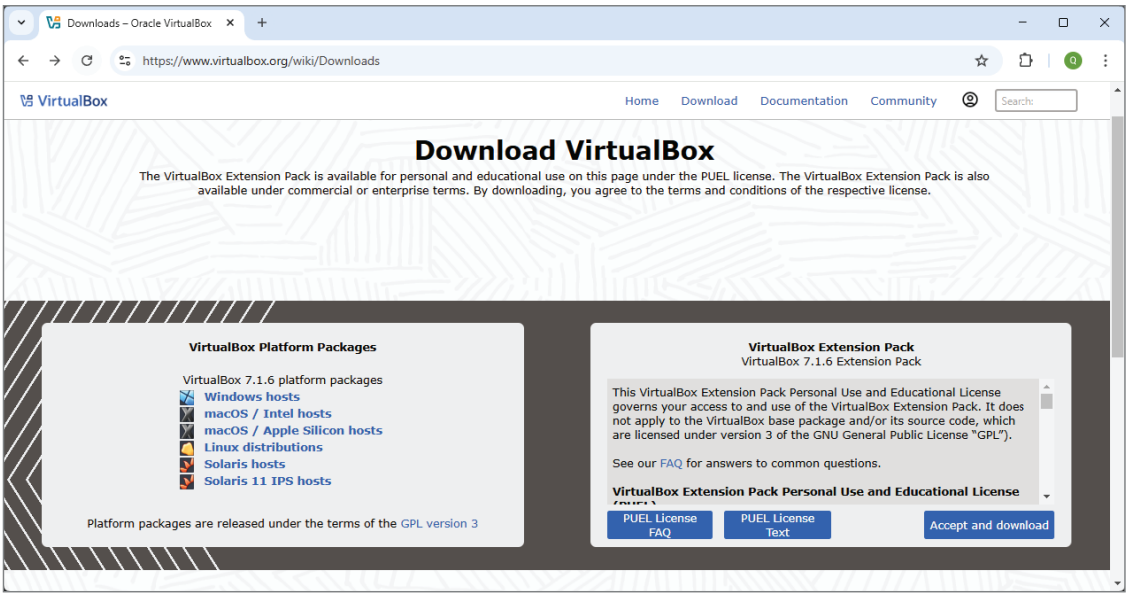
Downloads – Oracle VirtualBox
https://www.virtualbox.org/wiki/Downloads
VirtualBox
Home
Download
Documentation
Community
Search
Download VirtualBox
The VirtualBox Extension Pack is available for personal and educational use on this page under the PUEL license. The VirtualBox Extension Pack is also available under commercial or enterprise terms. By downloading, you agree to the terms and conditions of the respective license.
VirtualBox Platform Packages
VirtualBox 7.1.6 platform packages
Windows hosts
macOS / Intel hosts
macOS / Apple Silicon hosts
Linux distributions
Solaris hosts
Solaris 11 IPS hosts
Platform packages are released under the terms of the GPL version 3
VirtualBox Extension Pack
VirtualBox 7.1.6 Extension Pack
This VirtualBox Extension Pack Personal Use and Educational License governs your access to and use of the VirtualBox Extension Pack. It does not apply to the VirtualBox base package and/or its source code, which are licensed under version 3 of the GNU General Public License "GPL").
See our FAQ for answers to common questions.
VirtualBox Extension Pack Personal Use and Educational License
PUEL License FAQ
PUEL License Text
Accept and download

EXERCISE 11.2 *(continued)*

2. Download Lubuntu from `https://lubuntu.me/downloads`. Download the current release by clicking the button for Desktop 64-bit. As of this writing, the latest stable release is 24.10. It will download a zipped file with an .iso extension. You will need that ISO file later; it will essentially act as a bootable CD for your OS installation.

Installing VirtualBox

Now you can begin the installation of VirtualBox.

3. Double-click the VirtualBox icon. If you get a security warning, click the Run button. Then click Next on the Setup Wizard screen.
4. On the Custom Setup screen, click Next and then Next again. It will give you a warning about your network interfaces. Click Yes. (Your network connections will come back automatically.)
5. Click Install.

   The installation may take several minutes. You may also need to clear a Windows User Account Control security pop-up warning box by clicking Yes to allow Oracle Corporation to install a device driver. There may also be another Windows Security box to clear as well by clicking Install.
6. Once the installation is complete, click Finish.

Configuring VirtualBox

It's time to configure VirtualBox.

7. You might get a VirtualBox warning telling you that an image file is not currently accessible. That's fine. Click Ignore. You should see a screen similar to the one shown in Figure 11.8.

FIGURE 11.8 VirtualBox preconfiguration

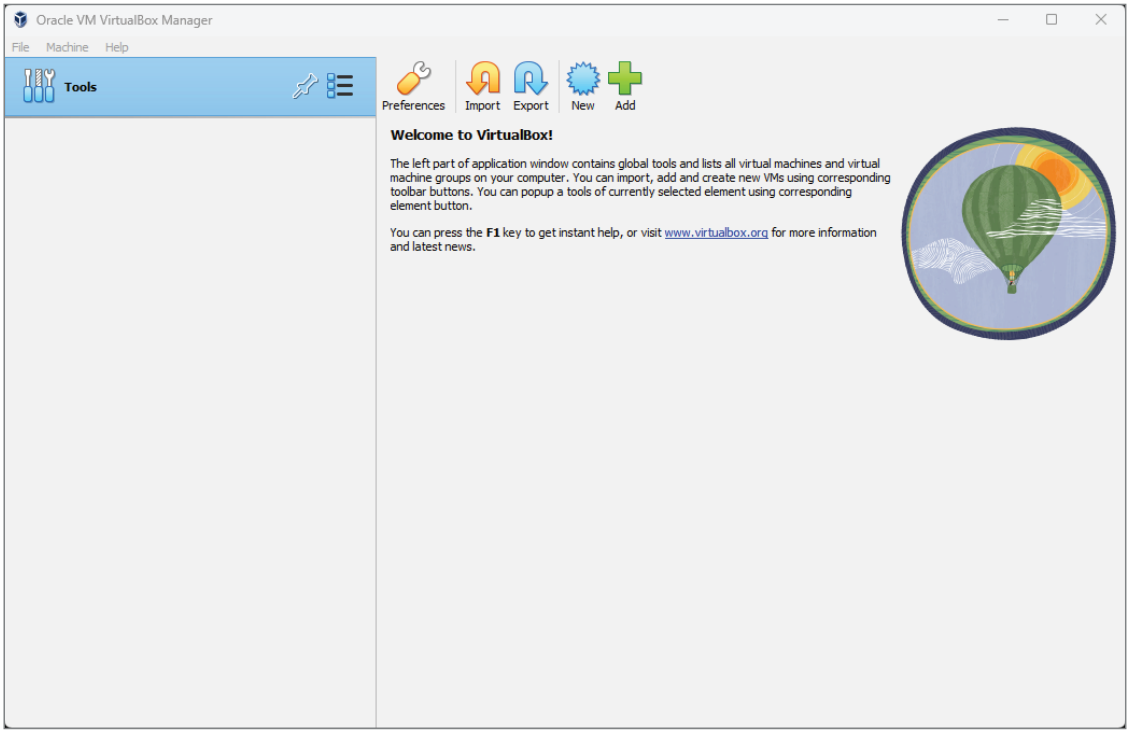
Oracle VM VirtualBox Manager
File Machine Help
Tools
Preferences Import Export New Add
Welcome to VirtualBox!
The left part of application window contains global tools and lists all virtual machines and virtual machine groups on your computer. You can import, add and create new VMs using corresponding toolbar buttons. You can popup a tools of currently selected element using corresponding element button.
You can press the F1 key to get instant help, or visit www.virtualbox.org for more information and latest news.

8. Click the blue New icon to create a new virtual machine. Give it a name. The Type and Version boxes aren't critical; they don't affect anything. If you type in **Lubuntu** for a name, it will automatically set Type to Linux and Version to Ubuntu. Click Next.
9. In the Hardware window, click Next to accept the default base memory and processor allocation.
10. In the Virtual Hard Drive window, the default option is Create A Virtual Hard Disk Now. Leave that option selected and click Next. (If you were going to use this extensively, you would want to give the virtual machine more hard drive space than the default 10 GB.)
11. The Summary page appears with your choices. Click Finish. You should get a screen similar to the one shown in Figure 11.9.

EXERCISE 11.2 *(continued)*

FIGURE 11.9 VirtualBox with a virtual drive

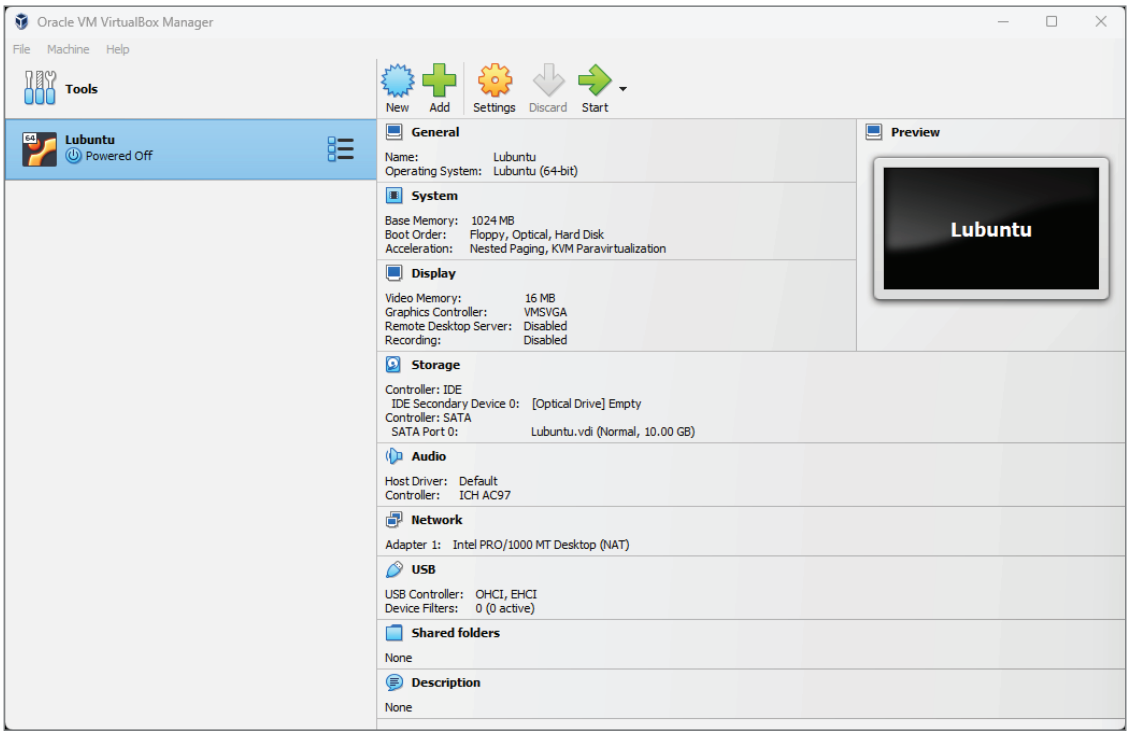

Great! You now have a virtual machine on your hard drive.

Installing the OS

Now you just need to put something on your VM—more specifically, an OS.

12. Click the Settings button.
13. In the Lubuntu – Settings window, click the Storage icon on the left.
14. Under one of your controllers, you should see something that looks like a disc icon that says Empty. It should look like Figure 11.10.

FIGURE 11.10 Lubuntu – Settings window

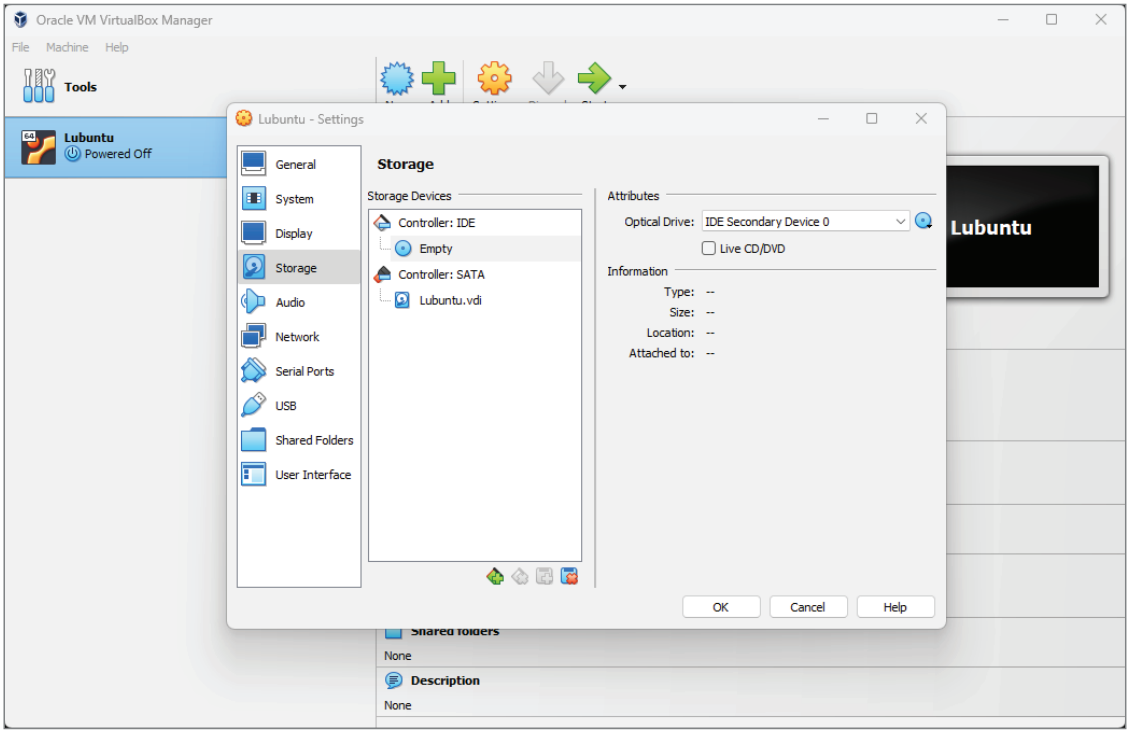

15. On the very right side of the window, you will have another disc icon with a little down arrow on it. Click that, and a menu will pop up. Select Choose/Create A Virtual Optical Disk.
16. The Optical Disk Selector window will appear. In this case, the Lubuntu ISO file is listed. If yours is not, you will need to navigate to the directory where you stored the `Lubuntu.iso` file that you downloaded. Highlight the file, and then click Choose.
17. Back on the Lubuntu – Settings screen, your drive that was empty should now say Lubuntu. Click OK.
18. Now you are back to the Oracle VirtualBox Manager screen. With Lubuntu on the left highlighted, click the green Start arrow. This will boot into a live session of Lubuntu, like what you see in Figure 11.11.

In step 18 after you press the green Start arrow, you might get an error message saying the operating system was not bootable. If that happens, use the drop-down box at the bottom of the error message to point to the ISO file again, and retry. It should work. If not, it's possible the ISO file is corrupted. You can download again or try downloading a different distribution to make it work, starting back at step 12.

EXERCISE 11.2 *(continued)*

FIGURE 11.11 Lubuntu live session

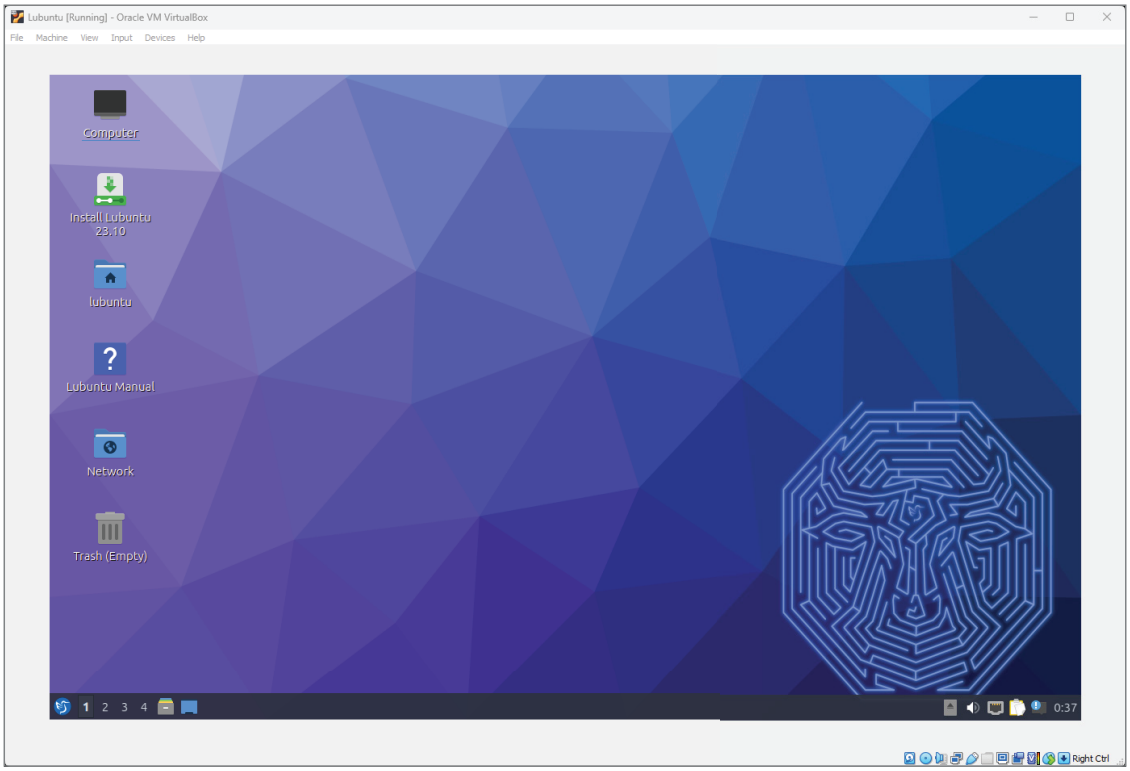

19. Double-click the Install Lubuntu icon on the virtual desktop to begin installation.
20. Choose a language and click Next.
21. Choose a location and click Next.
22. Select a keyboard layout and click Next.
23. You will get to a screen similar to the one shown in Figure 11.12, which asks you for an installation type. It looks scary, but choose the Erase Disk option. This will install it on the virtual disk that you created earlier with VirtualBox, and it will *not* wipe out your entire hard drive. Click Next.

FIGURE 11.12 Installing Lubuntu

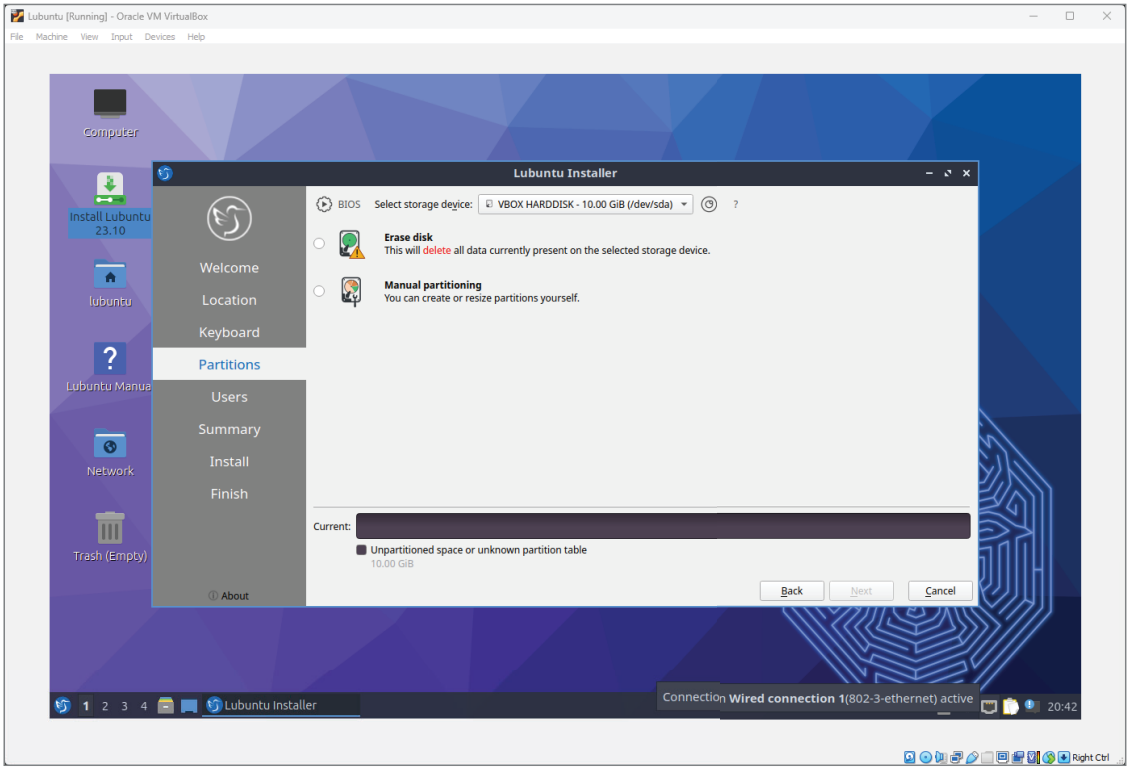

24. Continue with the installation process. When in doubt, choose the default and move to the next step.

25. When the installation is complete, select Restart Now and click Done.

Now that the installation is complete, play around in your new operating system! Here are some suggested exercises:

- Creating a user account
- Managing storage space
- Investigating file attributes
- Creating, moving, copying, and deleting files
- Creating a desktop shortcut
- Changing the desktop background and screen resolution

Understanding Cloud Computing

You hear the term a lot today—*the cloud*. What exactly is the cloud? The way it's named, and it's probably due to the word *the* at the beginning, it sounds like it's one giant, fluffy, magical entity that does everything you could ever want a computer to do. Only it's not quite that big, fluffy, or magical, and it's not even one thing.

Cloud computing is a method by which you access remote servers to store files or run applications for you. There isn't just one cloud—there are hundreds of commercial clouds in existence today. Many of them are owned by big companies, such as Microsoft, Google, HP, Apple, and Amazon. Basically, they set up the hardware and/or software for you on their network, and then you use it over the Internet.

Using the cloud sounds pretty simple, and in most cases it is no harder than using a web browser or Windows File Explorer. From the administrator's side, though, things can be a little trickier. Cloud computing relies upon virtualization, which as you learned earlier in this chapter means that there isn't necessarily a one-to-one relationship between a physical server and a logical (or virtual) server. In other words, there might be one physical server that virtually hosts cloud servers for a dozen companies, or there might be several physical servers working together as one logical cloud. From the end user's side, the idea of a physical machine versus a virtual machine doesn't even come into play, because it's all handled behind the scenes.

There are many advantages to cloud computing, and most of them revolve around money. Cloud providers can get economies of scale by having a big pool of resources available to share among many clients. It may be entirely possible for them to add more clients without needing to add new hardware, which results in greater profit. From a client company's standpoint, the company can pay for only the resources it needs without investing large amounts of capital into hardware that will be outdated in a few years. Using the cloud is often cheaper than the alternative. Plus, if there is a hardware failure within the cloud, the provider handles it. If the cloud is set up right, the client won't even know that a failure occurred. Other advantages of cloud computing include fast scalability for clients and ease of access to resources regardless of location.

The biggest downside of the cloud has been security. The client company's data is stored on someone else's server (off-premises), and company employees are sending it back and forth via the Internet. Cloud providers have dramatically increased their security over the last several years, but this can still be an issue, especially if the data is highly sensitive material or personally identifiable information (PII). Also, some companies don't like the fact that they don't own the physical assets.

Now let's dive into the types of services clouds provide, the types of clouds, cloud-specific terms with which you should be familiar, and some examples of using a cloud from the client side.

Types of Clouds

Most of the time when we think about clouds, they are run by big IT companies with tons of resources. This traditional type of cloud is called a *public cloud*. These clouds offer the

best in scalability, reliability, flexibility, geographical independence, and cost effectiveness. Whatever the client wants, the client gets. For example, if the client needs more resources, it simply scales up and uses more. Of course, the client will also pay more, but that's part of the deal.

Exam Objective 4.2, "Summarize cloud computing concepts," specifies that you should know about public, private, hybrid, and community clouds.

Setting up a cloud is not restricted to big companies offering services over the Internet, however. Any company can purchase virtualization software to set up its own individual cloud within its own network. That type of setup is referred to as a *private cloud*. Running a private cloud pretty much eliminates many of the features that companies want from the cloud, such as rapid scalability and eliminating the need to purchase and manage computer assets. The big advantage, though, is that it allows the company to control its own security within the cloud.

Some clients have chosen to combine public and private clouds into a *hybrid cloud*. This gives the client the great features of a public cloud while simultaneously allowing for the storage of more sensitive information on the private cloud. It's the best of both worlds.

The last type of cloud to discuss is a *community cloud*. These are created when multiple organizations with common interests combine to create a cloud. In a sense, it's like a public cloud but with better security. The clients know who the other clients are and, in theory, can trust them more than they could trust random people on the Internet. The economies of scale and flexibility won't be as great as with a public cloud, but that's the trade-off for better security.

Cloud Services

Cloud providers sell everything "as a service." The type of service is named for the highest level of technology provided. The basic levels, from, lowest to highest, are hardware (or infrastructure), software development tools (or the platform), and then software or applications themselves. For example, if computing and storage is the highest level purchased, the client will get infrastructure as a service. If applications are involved, it will be software as a service. Nearly everything that can be digitized can be provided as a service. Let's take a look at the three most common types of services offered by cloud providers, from the ground up:

Infrastructure as a Service Let's say that a company needs extra network capacity, including processing power, storage, and networking services (such as firewalls) but doesn't have the money to buy more network hardware. Instead, it can purchase *infrastructure as a service (IaaS)*, which is a lot like paying for utilities—the client pays for what it uses. Of the three major cloud services, IaaS requires the most network management expertise from the client. In an IaaS setup, the client provides and manages the software.

IaaS, PaaS, and SaaS are included in Exam Objective 4.2. Understand what each of them provides.

Platform as a Service *Platform as a service (PaaS)* adds a layer to IaaS that includes software development tools such as runtime environments. Because of this, it can be very helpful to software developers; the cloud vendor manages the various hardware platforms. This frees up the software developer to focus on building their application and scaling it. The best PaaS solutions allow for the client to export their developed programs and run them in an environment other than where they were developed. Examples of PaaS include Google App Engine, Microsoft Azure, Red Hat OpenShift, Amazon Web Services (AWS) Elastic Beanstalk, Engine Yard, and Heroku.

Software as a Service (SaaS) The highest of these three levels of service is *software as a service (SaaS)*, which handles the task of managing software and its deployment, and includes the platform and infrastructure as well. This is the one with which you are probably most familiar, because it's the model used by Google Docs, Microsoft 365, and even storage solutions such as Dropbox. The advantage of this model is to cut costs for software ownership and management; clients typically sign up for subscriptions to use the software and can renew as needed.

Figure 11.13 shows examples of these three types of services. SaaS is the same as the Application layer shown in the figure.

The three service levels included in the A+ exam objectives are the most common, but there are several other options you may see in the real world. The list here is by no means exhaustive, but should give you a good idea of the different types of services available from cloud providers:

- *Hardware as a service (HaaS)*, which is similar to IaaS but is more likely related specifically to data storage
- *Communications as a service (CaaS)*, which provides things like Voice over IP (VoIP), instant messaging, and video collaboration
- *Network as a service (NaaS)*, which provides network infrastructure
- *Desktop as a service (DaaS)*, which provides virtual desktops so that users with multiple devices or platforms can have a similar desktop experience across all systems
- *Data as a service (also DaaS)*, which provides for multiple sources of data in a mash-up
- *Business processes as a service (BpaaS)*, which provides business processes such as payroll, IT help desk, or other services
- *Anything/everything as a service (XaaS)*, which is a combination of the services already discussed

FIGURE 11.13 Common cloud service levels

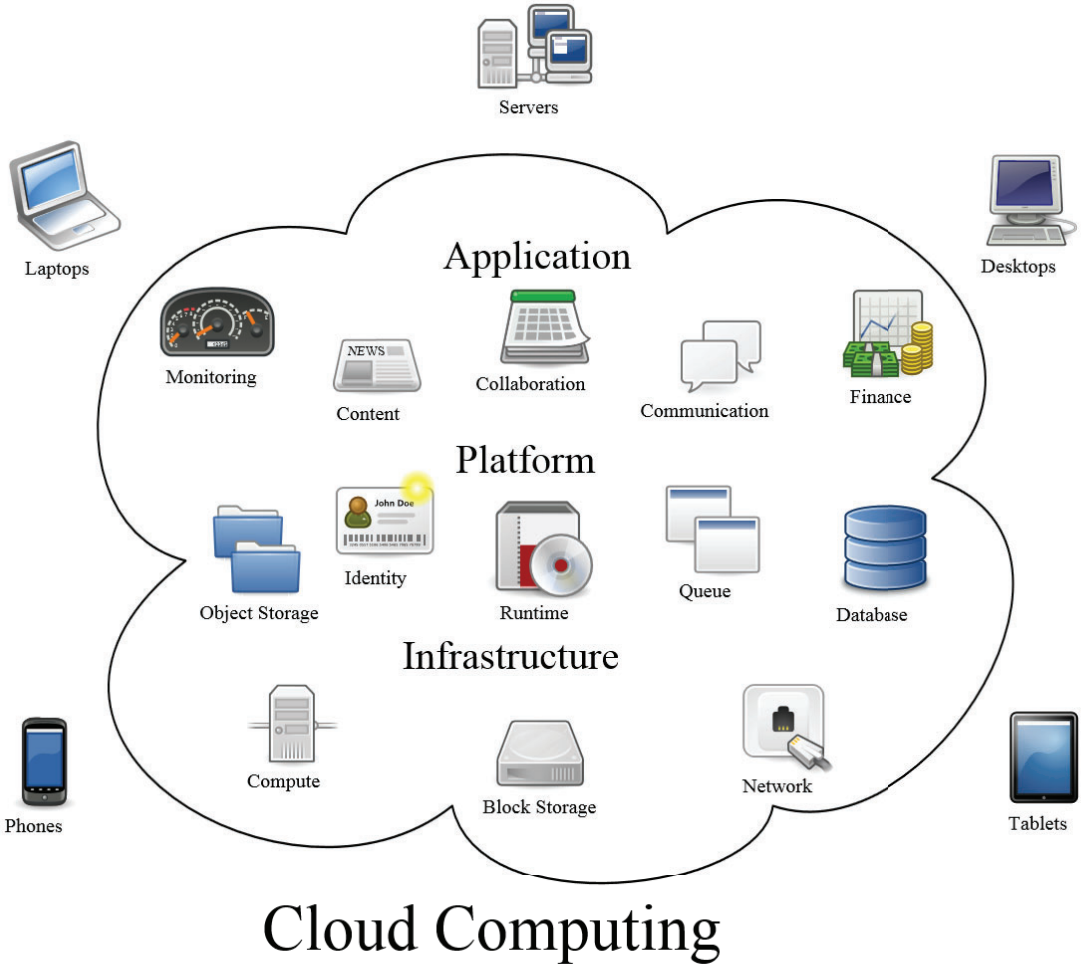

Source: Sam Johnston / Wikimedia Commons / https://commons.wikimedia.org/wiki/File:Cloud_computing.svg#/media/File:Cloud_computing.svg, last accessed on 07 March 2024 /CC BY 3.0.

The level of responsibility between the provider and the client is specified in the contract. It should be very clear which party has responsibility for specific elements, should anything go awry.

Important Cloud Characteristics

Up to this point, we've discussed types of clouds and levels of service. Users of the cloud will have expectations on the features of the cloud beyond who owns it or what service level is provided. For example, questions may include:

- Which types of devices can access the cloud?

- How does a client get more storage space and how long does it take to get it?
- How much uptime is guaranteed by the service provider?

To help clarify the characteristics that cloud providers should offer, the National Institute of Standards and Technology (NIST), a group within the U.S. Department of Commerce, has defined the following five essential characteristics of cloud computing:

On-Demand Self-Service This is one of the cloud's best features from an end user's standpoint. With on-demand self-service, users can access additional storage, processing, and capabilities automatically, without requiring intervention from the service provider.

Broad Network Access This means that cloud capabilities are accessible over the network by different types of clients, such as workstations, laptops, and mobile phones, using common access software such as web browsers. The ability for users to get the data they want, when they want, and how they want is sometimes referred to as *ubiquitous access*.

Resource Pooling, Shared Resources, or Multitenancy The idea of *resource pooling* is closely linked with virtualization, which we covered earlier in the chapter. It's essentially the same as the concept of *shared resources* and *multitenancy* (which are explicitly called out in the A+ exam objectives). Shared resources means that the provider's resources are seen as one large pool that can be divided up among clients as needed, and each client pays for the fraction of those resources they use. Multitenancy is almost the same thing—multiple clients (tenants) can use the same pool of hardware resources but keep their data and applications separate from each other.

Clients should be able to access additional resources as needed, even though the client may not be aware of where the resources are physically located. Typical pooled resources include network bandwidth, storage, processing power, and memory.

Public, hybrid, and community clouds use resource pooling or *shared resources*. In a private cloud, one company controls all the resources, which are called *dedicated resources*.

Rapid Elasticity We've talked about the ability to scale up resources as needed, and that is *elasticity*—a term used in the A+ exam objectives. In most cases, clients can get more resources instantly (or at least very quickly), and that is called *rapid elasticity*. For the client, this is a great feature because they can scale up without needing to purchase, install, and configure new hardware. Elasticity can also work backward; if fewer resources are required, the client may be able to scale down and pay less without needing to sell hardware. You will hear some subscriptions with built-in elasticity referred to as *pay-as-you-grow services*.

Measured Service Most cloud providers track clients' usage and then charge them for the services used. This type of setup is called *metered utilization* or *measured service*. Resource usage is monitored by the provider and reported to the client in a transparent fashion.

Within metered utilization there may be specific ingress or egress charges. *Ingress* refers to uploading data to the cloud, and with most cloud providers that service is free (of course, you have to pay for the storage space). *Egress* is the opposite of ingress—downloading data from the cloud. Many cloud providers will charge a fee for data egress.

Exam Objective 4.2 lists several cloud characteristics that you should be able to summarize for the exam. They include shared resources versus dedicated resources, metered utilization (and ingress/egress), elasticity, availability, file synchronization, and multitenancy.

In addition to those characteristics, A+ Exam Objective 4.2 lists two more you should know:

Availability When you contract for cloud services, it's reasonable to expect the services to be available when you need them. In the service level agreement (SLA), the cloud service provider (CSP) will be held accountable for ensuring a specific amount of uptime, which is known as *availability*. Sometimes you will hear it referred to as *high availability*, because of course higher is better.

Service availability is measured in terms of "nines," or how many nines of uptime the provider guarantees. For example, "three nines" means that the service will be available 99.9 percent of the time, whereas "four nines" will be up 99.99 percent of the time. More nines means more money, and different aspects of your service contract might require different levels of uptime. For example, a critical medical records database might need more guaranteed uptime than would a word processing application. The level of service you should get depends on how much risk your company is willing to take on and the trade-off with cost. Table 11.2 shows how much downtime is acceptable based on the number of nines of guaranteed uptime.

TABLE 11.2 Availability downtime

Availability	Downtime per year	Downtime per day
Three nines (99.9%)	8.77 hours	1.44 minutes
Four nines (99.99%)	52.6 minutes	8.64 seconds
Five nines (99.999%)	5.26 minutes	864 milliseconds
Six nines (99.9999%)	31.56 seconds	86.4 milliseconds

Guaranteeing that services will be available with the possible exception of less than one second per day seems pretty impressive, as is the case with five nines. You might see

other combinations, too, such as "four nines five," which translates into 99.995 percent availability, or no more than 4.32 seconds of downtime per day. The majority of CSPs will provide at least three nines or three nines five.

File Synchronization *File synchronization* is straightforward enough—it makes sure that the most current copy is on the cloud as well as on a local device. If changes are made to one, the other copy gets updated accordingly.

This section referred to the cloud characteristics as defined by the U.S. government agency tasked with promoting technical innovation. It's important to understand these characteristics because that's how cloud providers often refer to their services. Keep in mind that the A+ exam objectives specify a few different services and may use slightly different terms for the same services. We don't share both to confuse you, but rather to make sure you're familiar with industry terminology. For example, if one provider offers you measured service and another offers metered utilization, you should be able to recognize that they're the same thing.

Using Cloud-Based Services

Now that you understand cloud fundamentals such as models, service levels, and characteristics, let's turn our attention to some practical examples with which users will probably be more familiar. The two types of cloud interaction we will cover are storage and applications. The next two sections will assume the use of public clouds and standard web browser access.

Cloud-Based Storage

We introduced cloud storage in Volume 1 Chapter 10, "Mobile Connectivity and Application Support," because it's so integrated with the mobile user experience. It makes sense to review it quickly here as well, with a few key points to remember:

- A key purpose is to allow users access to their personal or work files regardless of their physical location or device they are using.
- Another great benefit is that teams of users, even geographically dispersed ones, can access the same files to aid collaboration.
- Pick a cloud provider that makes the most sense for the use case. If the entire corporate network is cloud-based, then ensure there is ample storage. If cloud storage is the only service needed, then find the cloud provider and plan with the right amount of space at an acceptable cost.
- Consider if other features are needed as well, such as apps to go along with the storage space (Google Workspace, for example).

Once the cloud-based storage is set up, ensure that all users are trained on how to store files to the cloud (it should be automatic or seamless to them) as well as access needed data and share nicely with others.

Cloud-Based Applications

Google really popularized the use of web-based applications. After all, the whole Chromebook platform, which has been very successful, is based on this premise. Other companies have gotten into the cloud-based application space as well, such as Microsoft with Microsoft 365. The cloud-based menus and layout are slightly different from PC-based versions, but if you're familiar with Word or Excel, you can easily use Microsoft 365—and all of the files are stored on the cloud.

Cloud-based apps run through your web browser. This is great for end users for a couple of reasons. One, your system does not have to use its own hardware to run the application; you are basically streaming a *virtual application*. Two, different client OSs can run the application (usually) without worrying about compatibility issues. Applications can often work across platforms as well, meaning that laptops, desktops, tablets, and smartphones can all use various apps.

Companies like Google have smartly integrated their apps into their cloud storage solutions. For example, you can access the Google Workspace apps from within Google Drive. Click or tap on the Google Apps icon (the nine dots in a square in the upper-right corner) to open the menu of available apps, as shown in Figure 11.14.

FIGURE 11.14 Google Apps

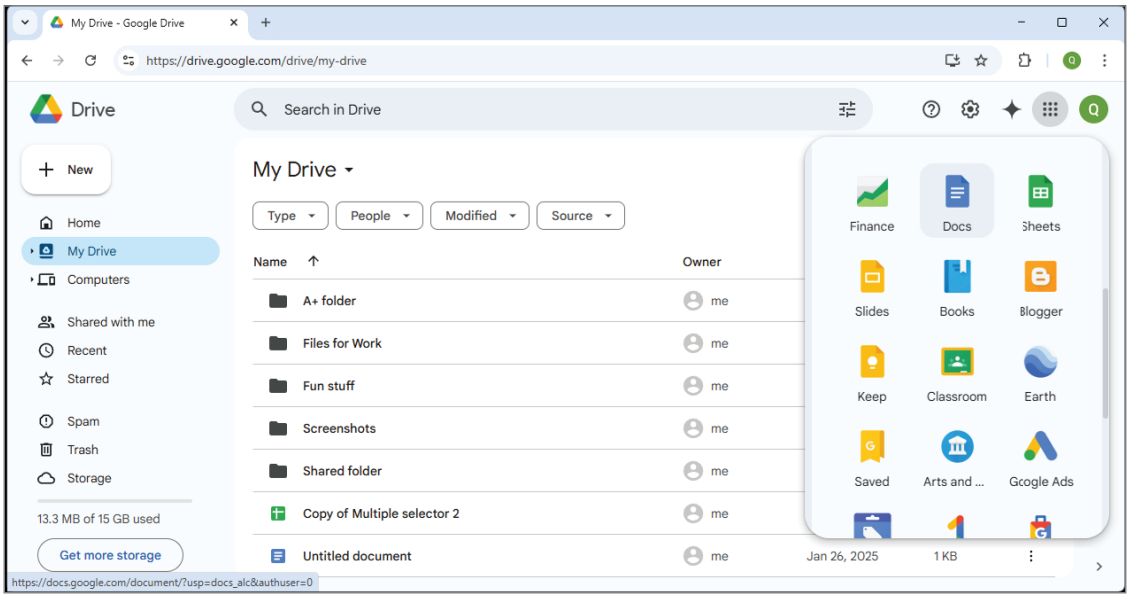

You can also create a new document using Google Docs by clicking the New button, as shown on the left side of Figure 11.14, and then choosing the application from the menu. If you choose Google Docs, it opens a new browser window with Google Docs, as shown in Figure 11.15. Notice that near the name of the document at the top, it says *Saved to Drive*. When it says this, you know that the document has been saved automatically.

FIGURE 11.15 Google Docs

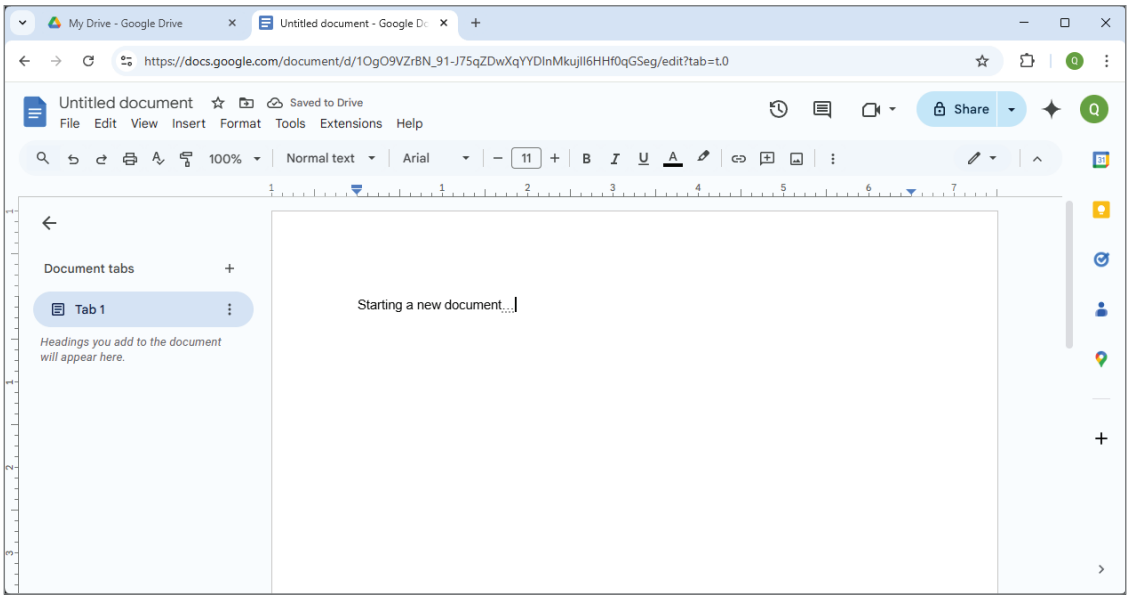

When selecting a cloud provider, you may use any one you like. In fact, it's better if you experience the differences in how providers store files and let you manage and manipulate them before making your choice. Exercise 11.3 will give you experience with using cloud-based storage and applications—specifically, Google Drive and its associated apps. This exercise will work best if you have someone you can work with. For example, in a classroom setting, you can partner with someone. If you are studying at home, you can create multiple accounts and get the same experience. You will just need to log out and in with your other account to see the shared files.

EXERCISE 11.3

Using Google's Cloud Services

1. Open Google at `www.google.com`.

   If you do not already have a Google account, you will need to create one. With it, you use Google's online apps and storage as well as a Gmail account.

2. If you are doing this exercise on your own, create a second account to share files and folders with.
3. Once you're logged in, click the Apps icon in the upper-right corner. Recall that the Google Apps icon is the nine dots, like you saw in Figure 11.14.
4. In Apps, click Drive. This will open Google Drive (refer to Figure 11.14).
5. Create a folder by clicking New ➢ New Folder.
6. Share the folder with another account by right-clicking the folder, then choosing Share ➢ Share (Figure 11.16).

FIGURE 11.16 Sharing a folder

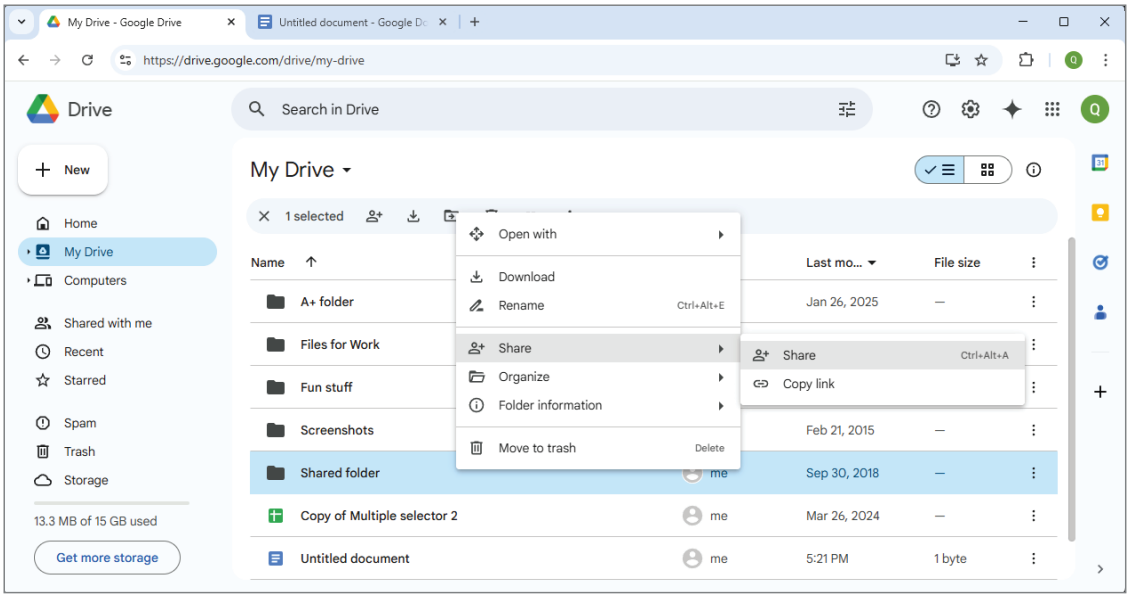

7. Also create a document or spreadsheet using Google's online software.

   How easy or difficult was it?
8. If necessary, log out and log in to the other account that you created to access the resources that were shared for you.

   How easy or difficult was it?

Another example of web applications and cloud storage is the streaming of media. Companies such as Netflix, Amazon, Spotify, Pandora, Apple, and others store movies and music on their clouds. You download their client software, and for a monthly subscription fee,

you can stream media to your device. It can be your phone, your tablet, your computer, or your home entertainment system. Before the advent of broadband network technologies, this type of setup would have been impossible, but now it has become the mainstream way that people receive audio and video entertainment.

Many companies have scrapped internal email management and moved to cloud-based email. By using an *off-site email application*, such as Gmail, a company can use their `.com` email address and Gmail as the client. This relieves administrative burden—all that's required is to call Google to set it up properly. There's a monthly fee, but it's often cheaper than managing an email server internally. And users can easily access their email from anywhere with an Internet connection. This is an example of SaaS because the cloud provider is hosting the email server and users access it using a browser, which is simply the client.

Setting Up a Cloud Server

In this chapter we've given you the foundational knowledge needed to set up public cloud services as a client. Basically, here are the steps:

1. Determine network and user needs.
2. Contact a big cloud service provider or two and get bids.
3. Sign up. The CSP will give you directions on how to access services and resources.
4. Train network users on how to use cloud services.

Some of you may be reading this chapter wondering how to set up a cloud server, though. (Or you might not, and that's okay too!) We will use this section to show you how. To be clear, setting up and managing a cloud server is far beyond the scope of what the A+ exam will ask you to do. Therefore, don't feel the need to memorize the material in this section. We include it because we think it's fun and might be interesting to you. For the cloud services in this section, we will use Amazon Web Services (AWS) Free Tier. It won't have all the functionality that a paid subscription does, but it has enough for us to play with and learn from. The first step is to create a free AWS account, which we show you how to do in Exercise 11.4.

EXERCISE 11.4

Creating a Free AWS Account

To use AWS, you need to have an account. AWS provides several types of accounts with different levels of service. Fortunately, users can create a free account for testing or learning purposes.

1. Go to `https://aws.amazon.com/free`.
2. Click the button that says Create a Free Account.

3. On the account creation page (Figure 11.17), enter the root user email address and AWS account name, then click Verify Email Address.

FIGURE 11.17 Creating AWS account

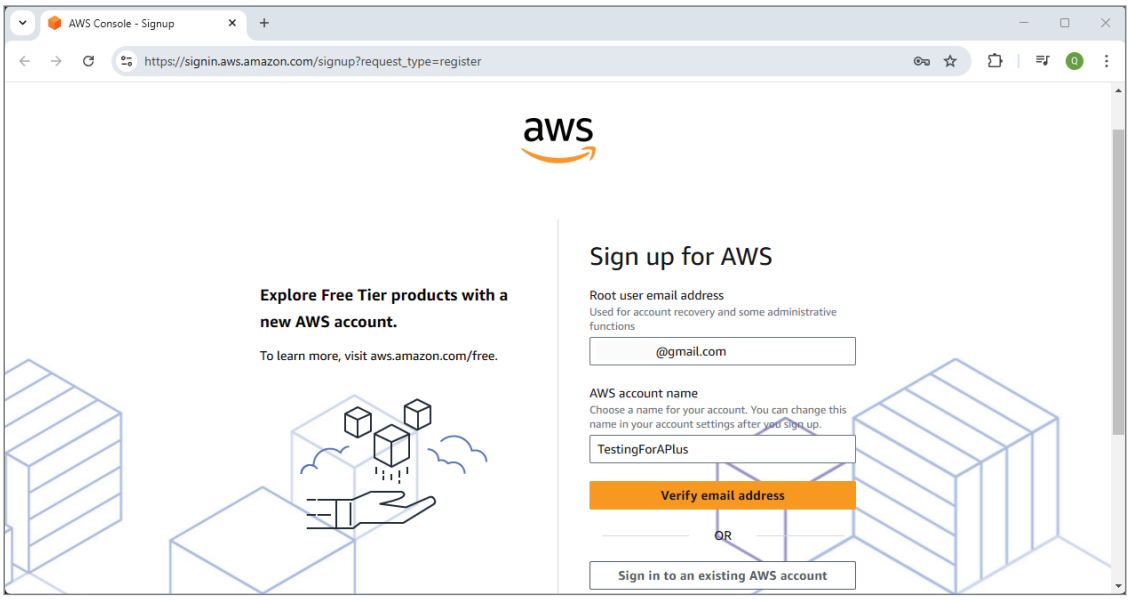
AWS Console - Signup
https://signin.aws.amazon.com/signup?request_type=register
aws
Explore Free Tier products with a new AWS account.
To learn more, visit aws.amazon.com/free.
Sign up for AWS
Root user email address
Used for account recovery and some administrative functions
@gmail.com
AWS account name
Choose a name for your account. You can change this name in your account settings after you sign up.
TestingForAPlus
Verify email address
OR
Sign in to an existing AWS account

4. Verify the code.
5. Create the root user password. Don't lose this!

   In the UNIX and Linux worlds, `root` is the all-powerful user with administrative access to everything. Its security information should be protected at all costs. (You will learn more about root and other Linux accounts in Volume 2 Chapter 4, "Working with macOS and Linux.")

6. Enter your contact information. For purposes of this exercise, choose a Personal account. Check the box that you have read the AWS Customer Agreement, and then click Agree And Continue.
7. Enter billing information. Amazon will not charge you for usage below its free limits. For what we will do in this chapter you should not incur any charges. Click Verify And Continue.
8. Confirm your identity. (Amazon is serious about security!)
9. Verify the code Amazon texted you.
10. Choose a support plan (Figure 11.18). Choose Basic Support - Free, unless you want one of the other options. Click Complete Sign Up.

EXERCISE 11.4 *(continued)*

FIGURE 11.18 Selecting a support plan

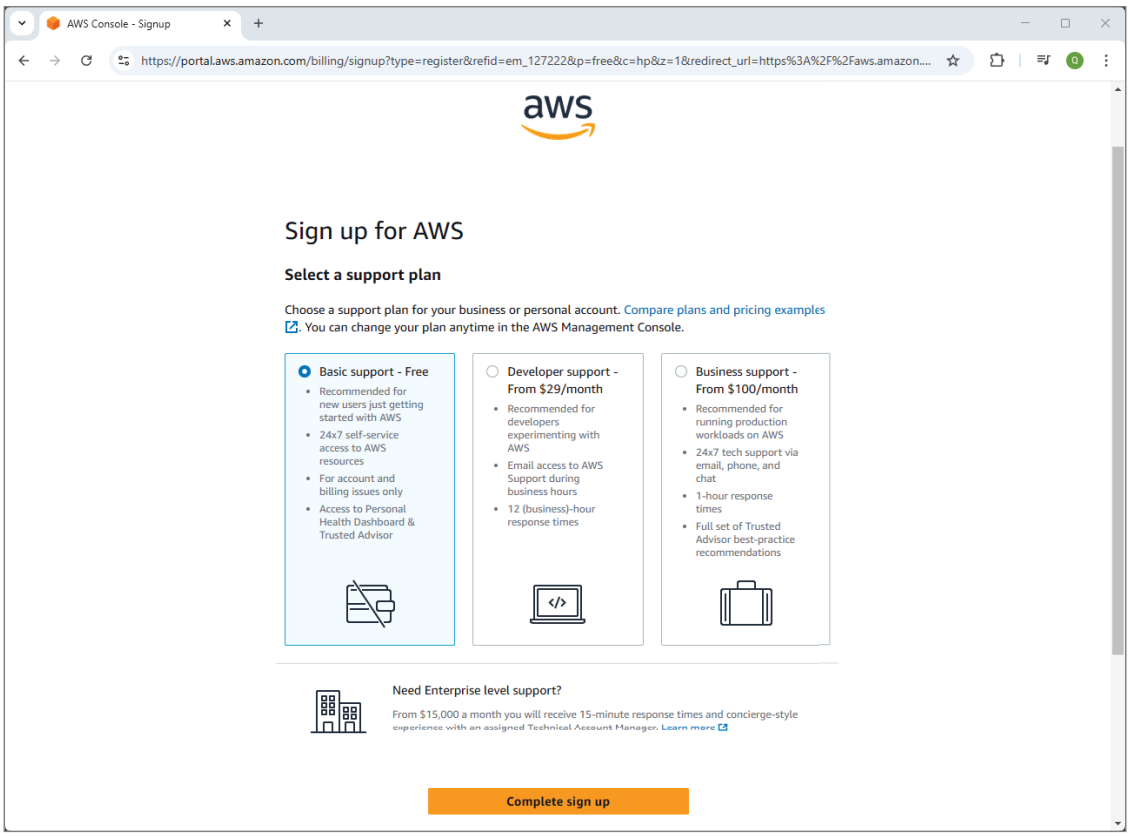
AWS Console - Signup
https://portal.aws.amazon.com/billing/signup?type=register&refid=em_127222&p=free&c=hp&z=1&redirect_url=https%3A%2F%2Faws.amazon....
aws
Sign up for AWS
Select a support plan
Choose a support plan for your business or personal account. Compare plans and pricing examples. You can change your plan anytime in the AWS Management Console.
Basic support - Free
Recommended for new users just getting started with AWS
24x7 self-service access to AWS resources
For account and billing issues only
Access to Personal Health Dashboard & Trusted Advisor
Developer support - From $29/month
Recommended for developers experimenting with AWS
Email access to AWS Support during business hours
12 (business)-hour response times
Business support - From $100/month
Recommended for running production workloads on AWS
24x7 tech support via email, phone, and chat
1-hour response times
Full set of Trusted Advisor best-practice recommendations
Need Enterprise level support?
From $15,000 a month you will receive 15-minute response times and concierge-style
Complete sign up

11. You will get a screen congratulating you, with a link to go to the AWS Management Console. Click the link if you'd like; it will take you to the page where we will start Exercise 11.5.

After creating the account, you have access to Amazon's free AWS services. Here are some of the common services you have access to:

- Elastic Compute Cloud (EC2), where you can create virtual machines
- Simple Storage Service (S3), which allows you to create and manage storage buckets
- Relational Database Service (RDS), which is a database management app
- Lambda, which is a service that will run code in response to an event, such as a user action or a failed login

Now that you have an account, Exercise 11.5 will take you through creating a virtual server.

EXERCISE 11.5

Creating a Virtual Server in AWS

In this exercise, you will create a virtual Windows server and learn how to access it for server management.

Creating the Server

1. Open the AWS Management Console. If you didn't click the link at the end of Exercise 11.4, you can go to `https://aws.amazon.com` and click Sign In To The Console to sign in.
2. The first screen you see when you log into the management console is Console Home (Figure 11.19). To create a virtual server, click the EC2 link in the Recently Visited section.

FIGURE 11.19 AWS Console Home

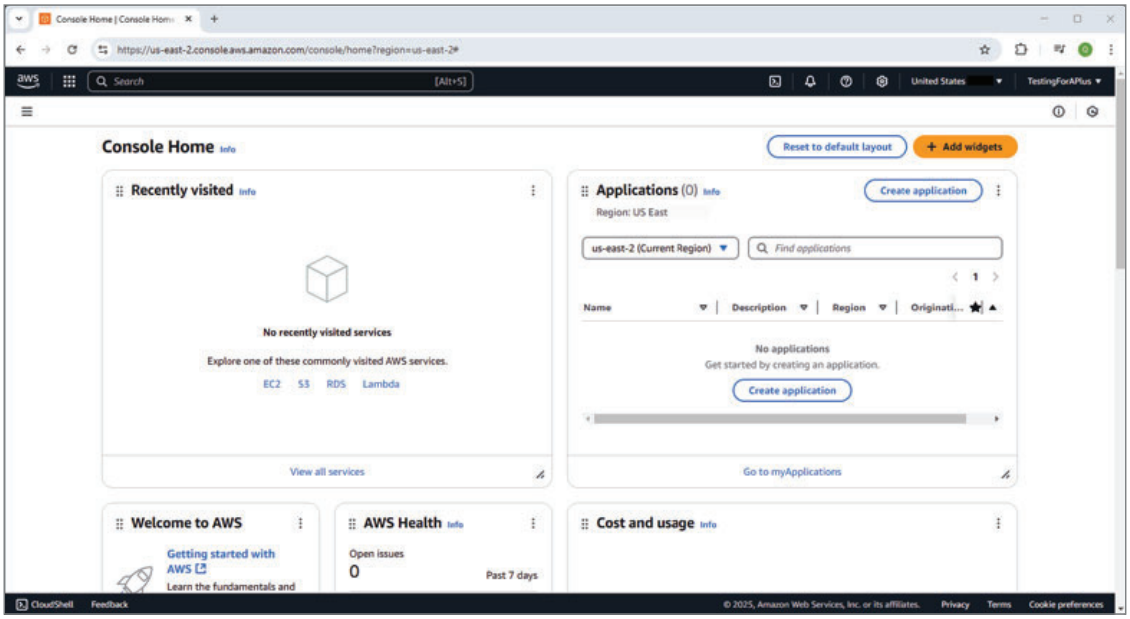

3. You will see an EC2 welcome screen. On the left-side navigation, click Dashboard (Figure 11.20). There's a lot to take in on this screen, but here's where you can see data about how many services or resources you are using, as well as set up virtual machines.

EXERCISE 11.5 *(continued)*

FIGURE 11.20 EC2 dashboard

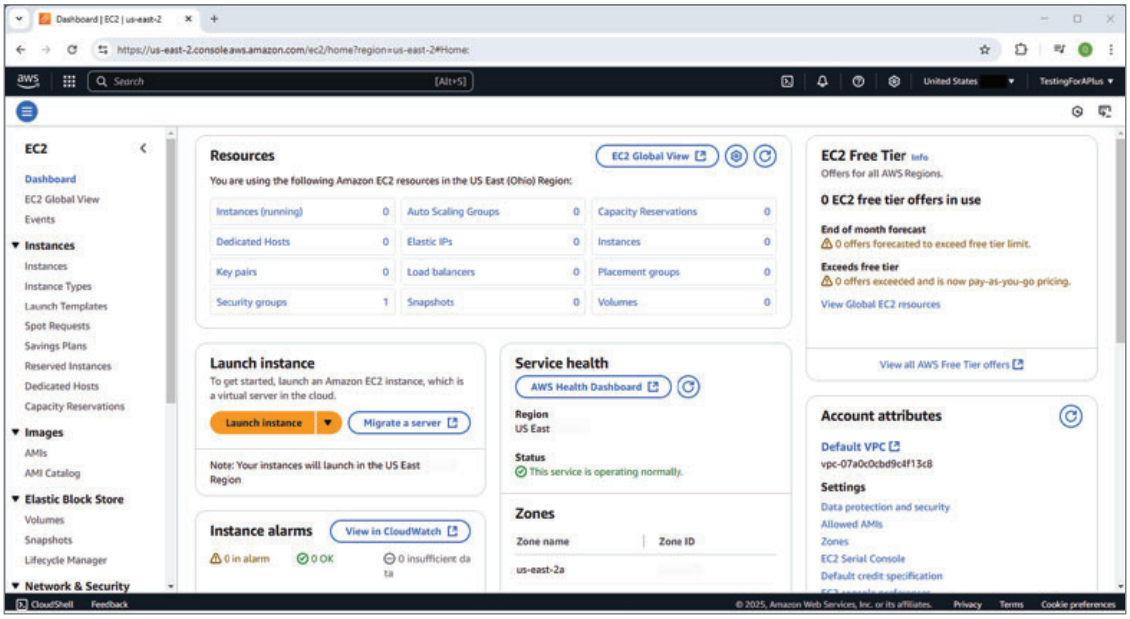

4. Click the orange button that says Launch Instance. You will get the Launch An Instance page (Figure 11.21).

FIGURE 11.21 Launch an Instance

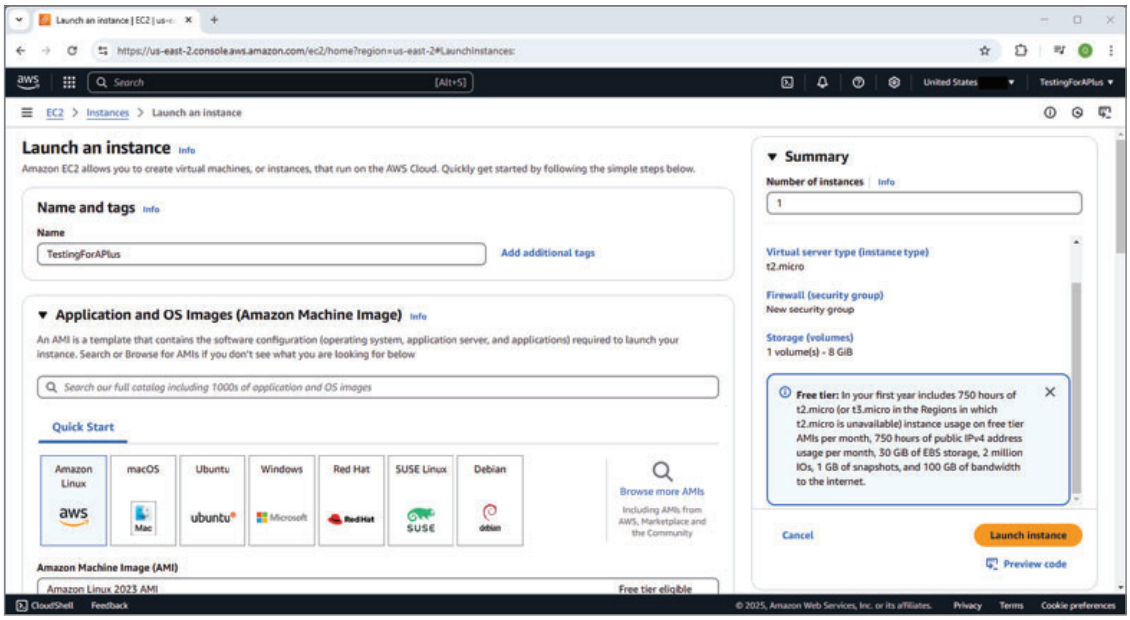
EC2 > Instances > Launch an instance
Launch an instance Info
Amazon EC2 allows you to create virtual machines, or instances, that run on the AWS Cloud. Quickly get started by following the simple steps below.
Name and tags Info
Name
TestingForAPlus
Add additional tags
Application and OS Images (Amazon Machine Image) Info
An AMI is a template that contains the software configuration (operating system, application server, and applications) required to launch your instance. Search or Browse for AMIs if you don't see what you are looking for below
Search our full catalog including 1000s of application and OS images
Quick Start
Amazon Linux
macOS
Ubuntu
Windows
Red Hat
SUSE Linux
Debian
Browse more AMIs
Including AMIs from AWS, Marketplace and the Community
Amazon Machine Image (AMI)
Amazon Linux 2023 AMI
Free tier eligible
Summary
Number of instances Info
1
Virtual server type (instance type)
t2.micro
Firewall (security group)
New security group
Storage (volumes)
1 volume(s) - 8 GiB
Free tier: In your first year includes 750 hours of t2.micro (or t3.micro in the Regions in which t2.micro is unavailable) instance usage on free tier AMIs per month, 750 hours of public IPv4 address usage per month, 30 GiB of EBS storage, 2 million IOs, 1 GB of snapshots, and 100 GB of bandwidth to the internet.
Cancel
Launch instance
Preview code
CloudShell
Feedback
Privacy
Terms
Cookie preferences

The left portion of the Launch an Instance page is where you make selections, and the right portion is the summary of choices you made. There are several options on this page. Notice that in the Application and OS Images section, there are several server OSs to choose from. This could be an opportunity for you later, if you'd like, to play around with different server OSs.

5. In the Application and OS Images section, click Windows. Notice that your selection changes what is available in the Amazon Machine Image section right below the OS selectors. As of this writing, the default option is Microsoft Windows Server 2025 Base. This is what we want (or some version of Microsoft Windows server).

6. Scroll down to the Instance Type section (Figure 11.22).

FIGURE 11.22 Instance Type

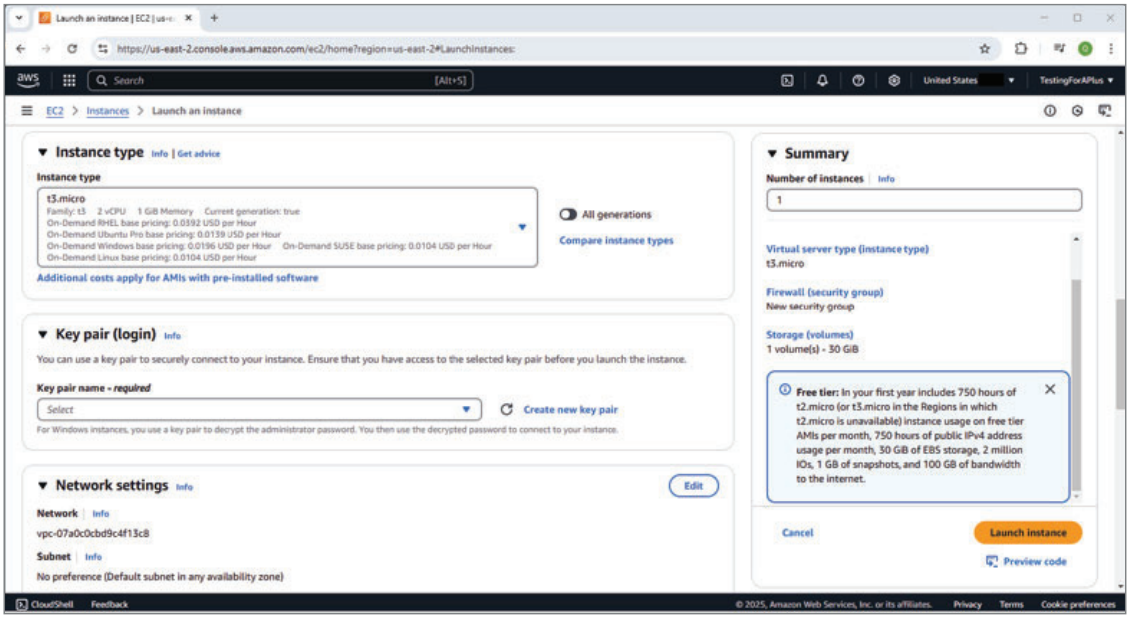
Launch an instance | EC2 | us-e
https://us-east-2.console.aws.amazon.com/ec2/home?region=us-east-2#LaunchInstances:
aws
Search
[Alt+S]
United States
TestingForAPIus
EC2 > Instances > Launch an instance
Instance type Info | Get advice
Instance type
t3.micro
Family: t3 2 vCPU 1 GiB Memory Current generation: true
On-Demand RHEL base pricing: 0.0392 USD per Hour
On-Demand Ubuntu Pro base pricing: 0.0139 USD per Hour
On-Demand Windows base pricing: 0.0196 USD per Hour On-Demand SUSE base pricing: 0.0104 USD per Hour
On-Demand Linux base pricing: 0.0104 USD per Hour
All generations
Compare instance types
Additional costs apply for AMIs with pre-installed software
Key pair (login) Info
You can use a key pair to securely connect to your instance. Ensure that you have access to the selected key pair before you launch the instance.
Key pair name - required
Select
Create new key pair
For Windows instances, you use a key pair to decrypt the administrator password. You then use the decrypted password to connect to your instance.
Network settings Info
Edit
Network Info
vpc-07a0c0cbd9c4f13c8
Subnet Info
No preference (Default subnet in any availability zone)
Summary
Number of instances Info
1
Virtual server type (instance type)
t3.micro
Firewall (security group)
New security group
Storage (volumes)
1 volume(s) - 30 GiB
Free tier: In your first year includes 750 hours of t2.micro (or t3.micro in the Regions in which t2.micro is unavailable) instance usage on free tier AMIs per month, 750 hours of public IPv4 address usage per month, 30 GB of EBS storage, 2 million IOs, 1 GB of snapshots, and 100 GB of bandwidth to the internet.
Cancel
Launch instance
Preview code
CloudShell
Feedback
© 2025, Amazon Web Services, Inc. or its affiliates.
Privacy
Terms
Cookie preferences

7. The default option is t3.micro, which is fine for our purposes. Notice how many virtual CPUs and memory are assigned. It's not much. Also notice how much various services cost. If you signed up for the free tier, you should not be charged unless you exceed the limits (and it tells you what those are on the right side of the screen). If you want, click the down arrow to the right of t3.micro to see what the different options are. Choose t3.micro.

8. The Key Pair section is very important. We need to create an encryption key pair to log in to the virtual server we're going to create. Click Create New Key Pair.

EXERCISE 11.5 *(continued)*

9. On the Create Key Pair screen (Figure 11.23), give the key pair a name. RSA is fine, as is the `.pem` file. Click Create Key Pair.

FIGURE 11.23 Create Key Pair

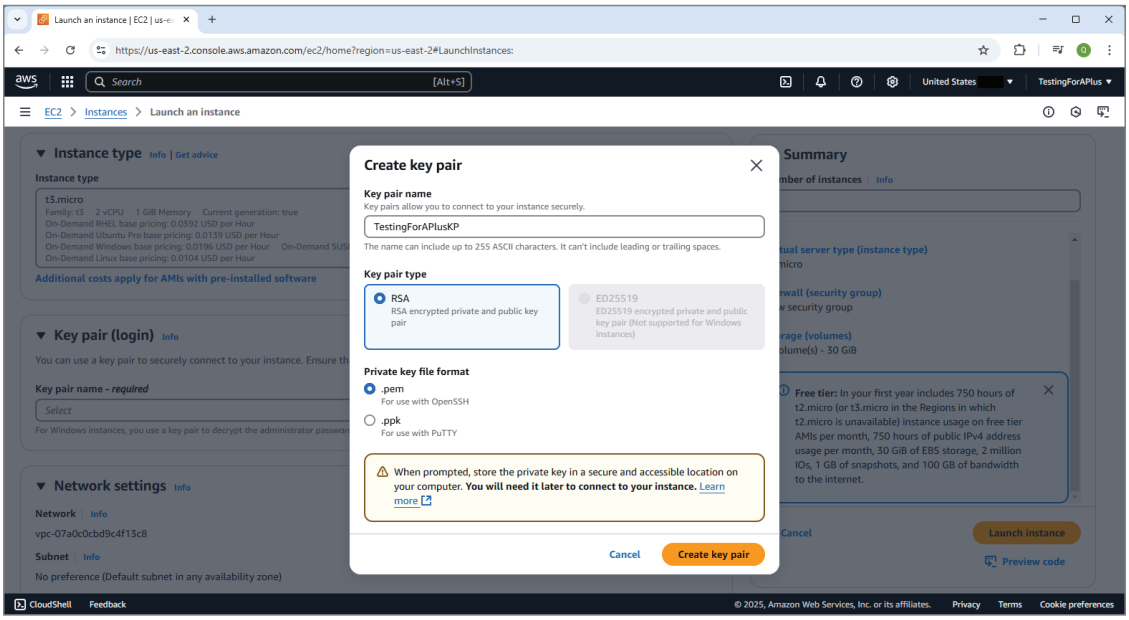
Launch an instance | EC2 | us-e
https://us-east-2.console.aws.amazon.com/ec2/home?region=us-east-2#LaunchInstances:
Search
[Alt+S]
United States
TestingForAPlus
EC2 > Instances > Launch an instance
Instance type
Key pair (login)
Network settings
Summary
Create key pair
Key pair name
Key pairs allow you to connect to your instance securely.
TestingForAPlusKP
The name can include up to 255 ASCII characters. It can't include leading or trailing spaces.
Key pair type
RSA
RSA encrypted private and public key pair
ED25519
ED25519 encrypted private and public key pair (Not supported for Windows instances)
Private key file format
.pem
For use with OpenSSH
.ppk
For use with PuTTY
When prompted, store the private key in a secure and accessible location on your computer. You will need it later to connect to your instance. Learn more
Cancel
Create key pair
Free tier: In your first year includes 750 hours of t2.micro (or t3.micro in the Regions in which t2.micro is unavailable) instance usage on free tier AMIs per month, 750 hours of public IPv4 address usage per month, 30 GiB of EBS storage, 2 million IOs, 1 GB of snapshots, and 100 GB of bandwidth to the internet.
Launch instance
Preview code
CloudShell
Feedback
© 2025, Amazon Web Services, Inc. or its affiliates.
Privacy
Terms
Cookie preferences

10. You will be asked to save the key pair `.pem` file. Choose a place where you will be able to easily find it.

    Do you remember the concept of public-private key pairs that we introduced in Volume 1 Chapter 8, "Network Services"? You are downloading your private key. This is just an exercise, but in real life you would never share your private key with anyone.

11. Scroll down to the Network Settings section. Look through the options. The defaults here are fine, with one change. Next to the Allow RDP Traffic From box, the default is Anywhere. For better security, click the dropdown arrow and choose My IP.

12. Scroll down to Configure Storage. This is where you could select additional storage space. For our purposes, the defaults are fine.

13. Scroll down to Advanced Details and click the drop-down twisty at the left side. There are dozens of options here. Read through them if you'd like. Again, for our purposes the defaults are fine.

14. On the right side, at the bottom of the Summary pane, click Launch Instance. It will take a moment, and then it should indicate Success at the top of the next screen. Scroll to the bottom and click View All Instances. That will take you to the EC2 Instances page (Figure 11.24). It should show the instance status as Running. The virtual server is alive!

FIGURE 11.24 EC2 Instances

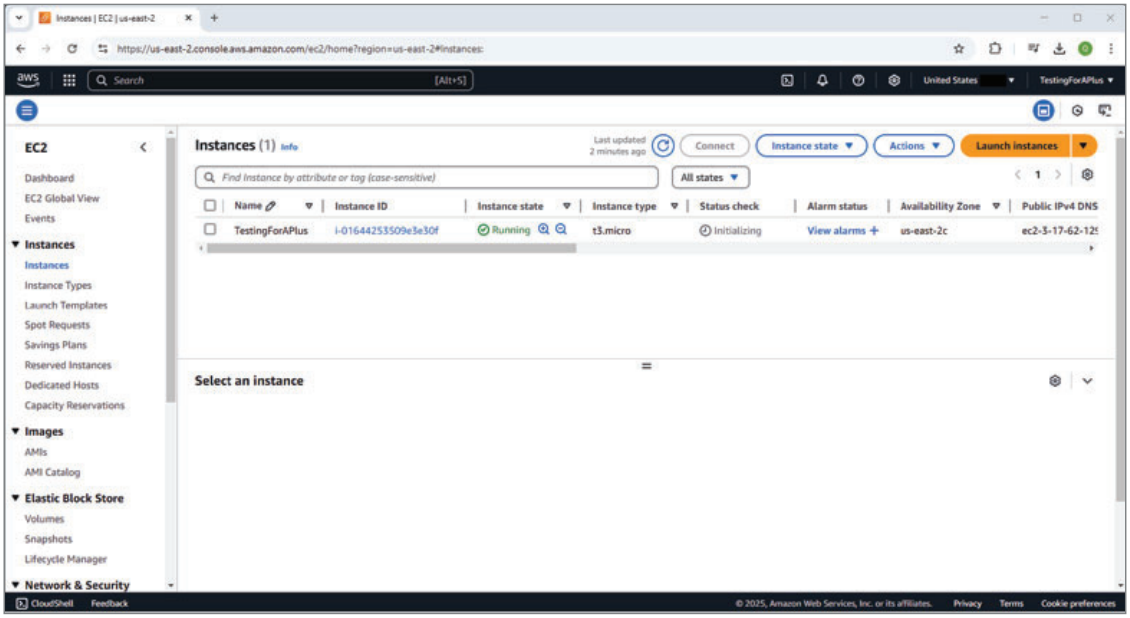

Examining the Virtual Server Instance

15. Click The Instance ID hyperlink shown in Figure 11.24. This will take you to the Instance Details shown in Figure 11.25.

FIGURE 11.25 Instance Details

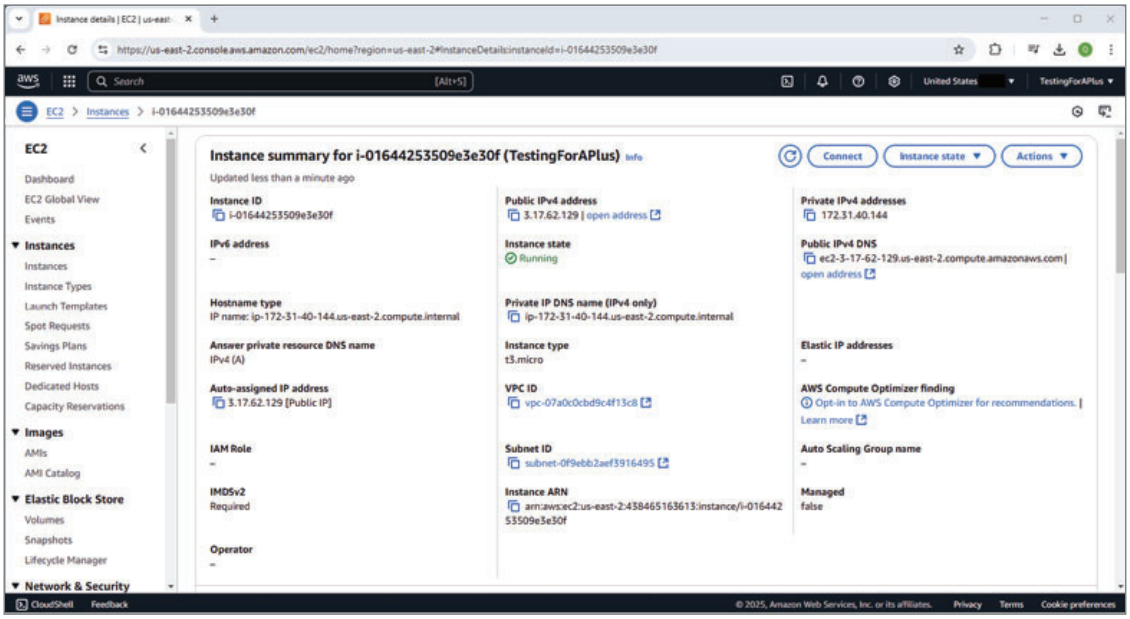

EXERCISE 11.5 *(continued)*

16. In the upper-right corner, click the Connect button. (There are also connect buttons on other pages, such as EC2 Instances, as shown in Figure 11.24). It will open the Connect To Instance page.

17. Click the RDP client tab to see a screen like the one in Figure 11.26.

FIGURE 11.26 Connect to Instance

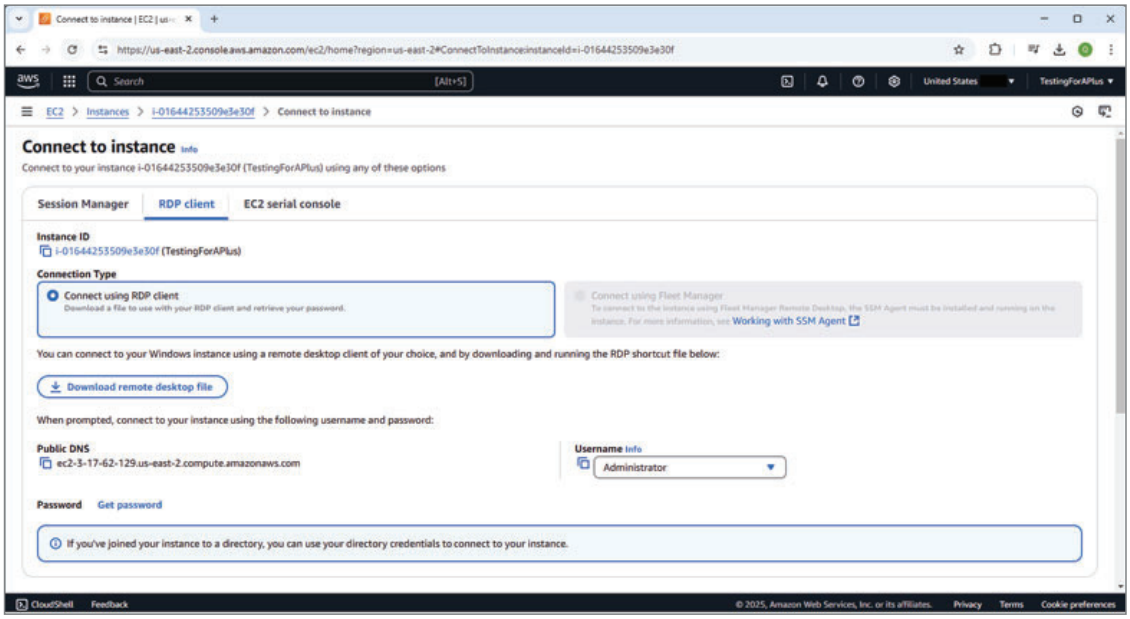
Connect to instance
Connect to your instance i-01644253509e3e30f (TestingForAPlus) using any of these options
Session Manager
RDP client
EC2 serial console
Instance ID
i-01644253509e3e30f (TestingForAPlus)
Connection Type
Connect using RDP client
Download a file to use with your RDP client and retrieve your password.
Connect using Fleet Manager
Working with SSM Agent
You can connect to your Windows instance using a remote desktop client of your choice, and by downloading and running the RDP shortcut file below:
Download remote desktop file
When prompted, connect to your instance using the following username and password:
Public DNS
ec2-3-17-62-129.us-east-2.compute.amazonaws.com
Username
Administrator
Password
Get password
If you've joined your instance to a directory, you can use your directory credentials to connect to your instance.
CloudShell
Feedback
Privacy
Terms
Cookie preferences

Do you remember remote desktop protocol (RDP) from Volume 1 Chapter 7, "Introduction to TCP/IP"? It's a real thing and you're going to get to use it!

18. Click on the Download Remote Desktop File button. Save the RDP client file to somewhere that's easy to find—the desktop is a good choice.

    You will use the RDP client file to access and manage the server using the RDP protocol. If you open the RDP client file now, it will ask you for the Windows Server login credentials. Do you remember setting an Administrator password for the Windows Server? No? It's probably because you didn't—AWS generated a long and complex password for you. It's time to get the password.

19. Click the Get Password link at the bottom of the page.

20. On the next screen (Get Windows Password), you will be asked to upload the private key you created in Step 10. Click the Upload Private Key File button, navigate to the private key, and select it. Once you have done that, the Private Key Contents window will be filled in. Click Decrypt Password.

    Now, the password will appear at the bottom of the of the RDP client page. Copy the password by using the icon to the left of it.

Accessing the Virtual Server

21. Find the RDP client you downloaded and open it. You may get a security warning. If so, click Connect.

22. In the Windows Security box (Figure 11.27), enter the password and click OK.

FIGURE 11.27 Windows Security

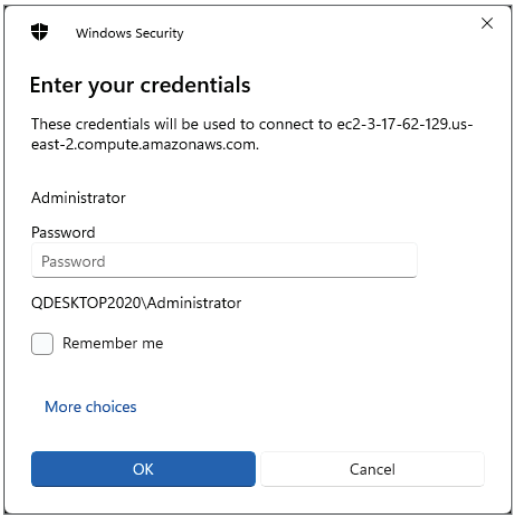

23. RDP may give you another security warning, stating that the identity of the remote server cannot be verified. Click Yes to continue. It will take a minute to connect, but the next screen you will see is a Windows Server 2025 interface (Figure 11.28).

FIGURE 11.28 Connect to Instance

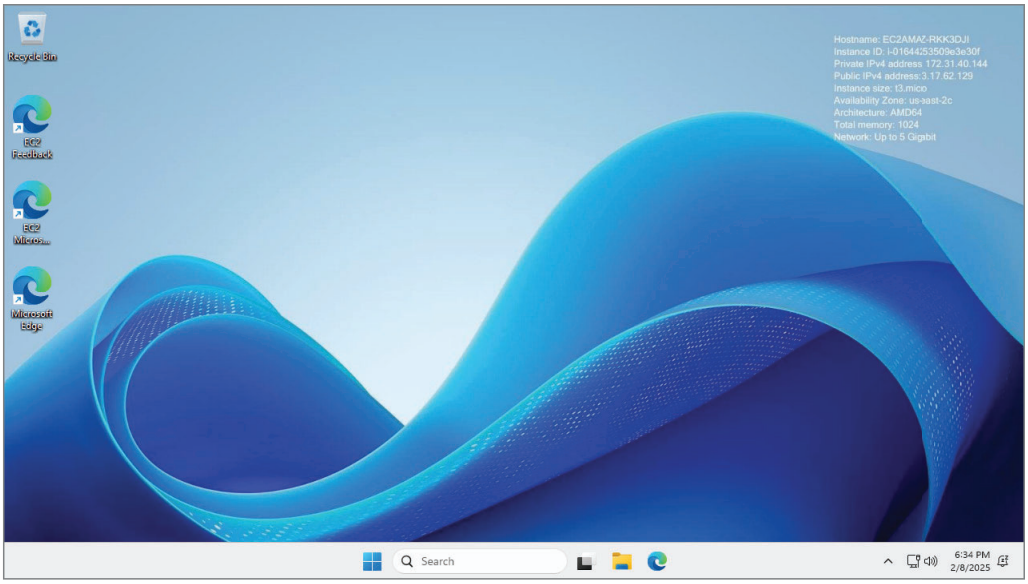

EXERCISE 11.5 *(continued)*

24. Play around in Windows server a bit. A few things to try include:
    - Connect to the Internet using Microsoft Edge.
    - Click Start and open File Explorer. Notice the folders are different than the computer you are on—you're connected to a completely separate computer in this RDP session.
    - Click Start ➢ Settings ➢ Accounts ➢ Other Users, and then Add Account to open the Local Users and Groups management console. Add a user to the server.
    - Click Start and open Server Manager. Play around a bit if you'd like to. For example, you could add DHCP or DNS services. (Remember, managing a Windows server is way beyond the scope of the A+ exam—this is just for fun.)

    How fast do tasks finish on the server? Probably a lot slower than you are used to. That's a feature of RDP—it will be slower than accessing apps on a local computer.
25. When you are done with your Windows server administration session, click the X on the right of the status bar at the top center of your RDP window.
26. You will get a warning about terminating your RDP session. Click OK.

Shutting Down the Virtual Server

27. In the AWS portal, navigate to EC2 Instances.
28. With the instance name highlighted, click the Instance State button above it.
29. Click Stop Instance to shut the server down. It will take a moment, and the Instance State will change to Stopped. If you want to access the server again, you will need to restart the instance.
30. If you're not going to use it again, terminate (delete) the instance. This is good virtual machine hygiene if the instance is no longer needed.

Summary

In this chapter, you learned about virtualization and cloud computing. Virtualization removes the barrier of there being a one-to-one relationship between computer hardware and an operating system. You learned the purpose of virtual machines, which includes sandboxing, test development, and application virtualization for legacy software and OSs and cross-platform virtualization, as well as requirements for desktop virtualization. You also learned about containers and the core piece of virtualization software, called the hypervisor.

The section finished with a long exercise on installing a hypervisor and Lubuntu on a Windows computer.

The next topic was cloud computing, which is virtualization in large scale over the Internet. Cloud computing has been one of the hottest topics in IT circles for several years now and will likely continue to be so for several more years. There are different types of clouds, such as public, private, community, and hybrid. Cloud providers sell several different types of services, such as IaaS, PaaS, and SaaS. Cloud features include shared resources, metered utilization, rapid elasticity, high availability, file synchronization, and multitenancy. You can use cloud services for storage, virtual applications (such as email or word processing), or both.

Exam Essentials

Understand the purpose of virtual machines. Virtual machines are designed to save providers and users money. They allow for multiple OSs to be installed on one computer. VMs can provide sandboxing, test development, and application virtualization for legacy software and OSs and cross-platform virtualization.

Understand virtual machine requirements. A virtual machine has security, network, and storage requirements. Those requirements will tie back to physical resources available on a server somewhere.

Understand what virtual desktops are. A virtual desktop is the collection of software and hardware needed to create a virtual environment. Sometimes it's called a virtual desktop infrastructure (VDI).

Know what the hypervisor is, and the differences between Type 1 and Type 2 hypervisors. The hypervisor is the software which allows the creation of virtual machines on a computer. A Type 1 (or bare-metal) hypervisor sits directly on the hardware and does not need a separate OS. This is mostly used for server-side virtualization. A Type 2 hypervisor sits on top of a host OS.

Know what containers are. Containers are an app and the required files (such as OS files and dependencies) needed to run the app in any environment.

Understand the four different types of clouds. A cloud can be public, private, hybrid, or community.

Know the differences between SaaS, IaaS, and PaaS. All of these are cloud terms. In infrastructure as a service (IaaS), the provider supplies the network infrastructure. In platform as a service (PaaS), software development tools and platforms are also provided. The highest level is software as a service (SaaS), where the provider supplies the infrastructure and applications.

Understand shared resources versus dedicated resources. All clouds use shared resources, which can be internal or external. A pool of resources is purchased and each participant in the cloud pays for part of it. The resources are shared resources. Dedicated resources would be wholly owned by one company.

Understand metered utilization, including ingress and egress. Metered utilization shows how much a client has used and will be billed for. Ingress is loading data to the cloud (which is usually free), and egress is downloading data from the cloud (which may come with an extra cost).

Know what elasticity and availability mean in the cloud. Elasticity means that a client can quickly get more (or fewer) resources as needed. Availability means that services are always or almost always available, such as three nines five.

Know what file synchronization means in the cloud. Synchronization means that an exact copy of the file is in multiple locations. Cloud file storage services include iCloud and Google Cloud, and in most cases have synchronization apps to sync to the mobile device.

Understand multitenancy. Multiple clients can use the same hardware resources from a single cloud provider, which is called multitenancy.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. You have been asked to identify the right type of cloud service to help the team of developers to provide programming elements such as runtime environments. Which service do you recommend?
   - **A.** PaaS
   - **B.** IaaS
   - **C.** SaaS
   - **D.** DaaS

2. You are setting up a cloud contract with a provider. Your team needs the ability to quickly increase capacity to meet peak demands. What do you request?
   - **A.** Elasticity
   - **B.** Availability
   - **C.** Resource pooling
   - **D.** Metered utilization

3. Your company wants to move to a cloud provider to be able to scale resources quickly, but it is concerned about the security of confidential information. Which of the following types of cloud models might be the most appropriate for your company?
   - **A.** Public
   - **B.** Private
   - **C.** Community
   - **D.** Hybrid

4. You have been asked to advise a group of several universities that want to combine research efforts and store data in the cloud. Which type of cloud solution might be best for them?
   - **A.** Public
   - **B.** Private
   - **C.** Community
   - **D.** Hybrid

5. Your manager wants to use the cloud because everyone seems to be talking about it. What should you include when you are listing the benefits of using the cloud? (Choose all that apply.)
   - **A.** Increased security
   - **B.** Increased scalability
   - **C.** Lower cost
   - **D.** Improved reliability

6. You have been asked to configure a client-side virtualization solution with three guest OSs. Each one needs Internet access. How should you configure the solution in the most cost-effective way?
   - **A.** Three physical NICs
   - **B.** One physical NIC, three virtual NICs, and one virtual switch
   - **C.** One physical NIC, one virtual NIC, and three virtual switches
   - **D.** One physical NIC, three virtual NICs, and three virtual switches

7. You have been asked to install Linux in a VM on a Windows 11 client. The Windows 11 client needs 4 GB of RAM, and Linux needs 2 GB of RAM. How much RAM does the system need at a minimum?
   - **A.** 4 GB
   - **B.** 6 GB
   - **C.** 8 GB
   - **D.** Unable to determine from the question

8. A computer using which of the following would be considered a legacy device? (Choose all that apply.)
   - **A.** A 386 processor
   - **B.** The IPX/SPX protocol
   - **C.** An application developed in 1989
   - **D.** 1 GB of RAM

9. You have been asked to set up client-side virtualization on an office computer. The host OS is Windows 11, and there will be three Windows 11 guest OSs. Which of the following is true about the need for antivirus security?
   - **A.** The host OS needs an antivirus program, but virtual machines can't be affected by viruses.
   - **B.** The host OS antivirus software will also protect the guest OSs on the VMs.
   - **C.** Installing antivirus software on the virtual switch will protect all guest OSs.
   - **D.** The host OS and each guest OS need their own antivirus software installed.

10. You need to set up a temporary operating system environment to quickly test a piece of software your manager wants to install on the network. Which should you use?
    - **A.** AAA server
    - **B.** Sandbox
    - **C.** SCADA
    - **D.** Application virtualization

11. You have been asked to set up client-side virtualization on a computer at work. The manager asks for a Type 2 hypervisor. What is the disadvantage of using that type of hypervisor?

A. The guest OS will compete for resources with the host OS.

B. The guest OS will be forced to a lower priortity with the CPU than the host OS.

C. The guest OS will be forced to use less RAM than the host OS.

D. The virtual guest OS will not be able to get on the physical network.

12. Your software development team has created an app that needs to be quickly deployed to all network users, using virtualization. Which technology is best suited for this purpose?

A. Type 1 hypervisor

B. Type 2 hypervisor

C. Containers

D. VDI

13. Which type of cloud is generally most cost effective for companies to use?

A. Private

B. Public

C. Hybrid

D. Community

14. A development team is creating containers. Which of the following statements are not true regarding containers? (Choose two.)

A. Containers are faster to deploy than hypervisors.

B. Each container has a minimum RAM requirement.

C. Containers can only hold one app.

D. Containers are designed to work on one OS.

15. Your manager read that multiple client companies might share the same hardware in a cloud and is extremely concerned about the security of company data. What should you tell them?

A. Data security is a serious concern in a multitenancy environment, so the data should be limited to non-sensitive materials.

B. Thanks to the use of containers, data from different companies is always separated to help avoid data security issues.

C. Cloud providers guarantee that clients will not share the same storge devices, so data security is not a significant issue.

D. Multiple tenants may share the same hardware, but their data and apps are separated to avoid data security issues.

16. A software developer wants to be able to test their in-progress apps in three different OSs. Which type of setup is best for their Windows 11 client computer?

A. Type 1 hypervisor

B. Type 2 hypervisor

C. Containers

D. VDI

17. A manager is demanding more storage space for his department, immediately. What type of cloud solution is best in this case?

A. IaaS

B. SaaS

C. PaaS

D. VDI

18. A cloud provider is asking if you know what all of the software and hardware is that's needed to create an appropriate virtual environment for your users. What are they referring to?

A. Type 1 hypervisor

B. Type 2 hypervisor

C. Containers

D. VDI

19. On the latest invoice from your CSP, you notice an egress surcharge. What is this most likely for?

A. Downloading data from the cloud.

B. Uploading data to the cloud.

C. Creating new containers for applications.

D. Creating new virtual desktop infrastructure (VDI) instances.

20. A CSP is setting up services for three client companies. Each client has contracted for 10 TB storage space. How much storage space does the cloud provider need to have, in total, to satisfy the contracts?

A. 10 TB, because clients will share hardware in a multitenancy environment.

B. 20 TB, because clients will share hardware in a multitenancy environment.

C. 30 TB, because that is what the clients contracted in total.

D. 60 TB, because the CSP must have reserve storage space in reserve for elasticity.

Performance-Based Question

You will encounter performance-based questions on the A+ exams. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answer compares to the authors', refer to Appendix B.

Describe the steps needed to enable Hyper-V in Windows 11.

Chapter

12

Core Hardware and Network Troubleshooting

THE FOLLOWING COMPTIA A+ 220-1201 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

- ✓ **2.0 Networking**
- ✓ **2.8 Explain networking tools and their purposes.**
  - Crimper
  - Cable stripper
  - Wi-Fi analyzer
  - Toner probe
  - Punchdown tool
  - Cable tester
  - Loopback plug
  - Network tap
- ✓ **5.0 Hardware and Network Troubleshooting**
- ✓ **5.1 Given a scenario, troubleshoot motherboards, RAM, CPUs, and power.**
  - Common symptoms
    - Power-on self-test (POST) beeps
    - Proprietary crash screens
    - Blank screen
    - No power
    - Sluggish performance
    - Overheating

- Burning smell
- Random shutdown
- Application crashes
- Unusual noise
- Capacitor swelling
- Inaccurate system date/time

✓ **5.2 Given a scenario, troubleshoot drive and RAID issues.**

- Common symptoms
  - Light-emitting diode (LED) status indicators
  - Grinding noises
  - Clicking sounds
  - Bootable device not found
  - Data loss/corruption
  - RAID failure
  - Self-monitoring, Analysis, and Reporting Technology (S.M.A.R.T.) failure
  - Extended read/write times
  - Low performance input/output operations per second (IOPS)
  - Missing drives in OS
  - Array missing
  - Audible alarms

✓ **5.5 Given a scenario, troubleshoot network issues.**

- Common symptoms
  - Intermittent wireless connectivity
  - Slow network speeds
  - Limited connectivity
  - Jitter
  - Poor Voice over Internet Protocol (VoIP) quality
  - Port flapping

- High latency
- External interference
- Authentication failures
- Intermittent Internet connectivity

Mentioning the words *troubleshooting theory* to many technicians can cause their eyes to roll back in their heads. It doesn't sound glamorous or sexy, and a lot of techs believe that the only way to solve a problem is just to dive right in and start working on it. Theories are for academics. In a way, they're right—you do need to dive in to solve problems because they don't just solve themselves. But to be successful at troubleshooting, you must take a systematic approach.

You may hear people say, "Troubleshooting is as much of an art as it is a science," and our personal favorite, "You just need to get more experience to be good at it." While there is an art to fixing problems, you can't ignore science. And if you need experience to be any good, why are some less experienced folks incredibly good at solving problems while their more seasoned counterparts seem to take forever to fix anything? More experience is good, but it's not a prerequisite to being a good troubleshooter. Again, it's all about applying a systematic approach.

There's one more detail to understand before getting into the details of specific problems: in order to troubleshoot anything, you need to have a base level of knowledge. For example, if you've never opened the hood of a car, it will be a bit challenging for you to figure out why your car won't start in the morning. If you're not a medical professional, you might not know why that body part hurts or how to make it feel better. In the same vein, if you don't know how data is stored and accessed on a computer, it's unlikely that you'll be able to fix related computer problems. So before you get too heavy into troubleshooting, make sure you understand how the systems on which you are working are supposed to function in the first place.

Because this chapter comes after the hardware and networking chapters, we're going to assume that you've read them already. Therefore, we're not going to get into a lot of detail about how things work—it's assumed that you know those details by now. (If you're still not certain, this book is a great reference manual!) Instead, we'll talk more about what happens when things don't work the way they're supposed to—what signs to look for and what to do to fix the problem.

After discussing some troubleshooting tricks and tips as well as a brief methodology, this chapter covers troubleshooting core hardware components, including motherboards, RAM, CPUs, and power, storage drive and RAID issues, and network issues. The remainder of the Hardware and Network Troubleshooting objective is covered in Volume 1 Chapter 13, "Mobile Device, Display, and Printer Troubleshooting."

Understanding Troubleshooting Best Practices

When troubleshooting, you should assess every problem systematically and try to isolate the root cause. Yes, there is a lot of art to troubleshooting, and experience plays a part too. But regardless of how "artful" or experienced you are, haphazard troubleshooting is doomed to fail. Conversely, even technicians with limited experience can be effective troubleshooters if they stick to the principles. The major key is to start with the issue and whittle away at it until you can narrow it down and pinpoint the problem.

No matter how skilled you may be at troubleshooting, always consider corporate policies, procedures, and impacts before implementing any changes. Don't perform any actions that would get you or the users into trouble, or that would otherwise cause harm to the company.

Troubleshooting Tricks and Tips

Throughout this chapter we will cover a variety of specific troubleshooting ideas, but there are some general tips and tricks that will serve you well regardless of the situation you're faced with. Although everyone approaches troubleshooting from a different perspective and with varying levels of knowledge, a few things should remain constant:

Checking the Easy Stuff First This might sound obvious, but in the sometimes frantic pace of trying to fix a problem, even the best techs occasionally forget this. The easy stuff means checking the connections or rebooting the computer. Everyone rolls their eyes when the tech asks if they rebooted, but a quick reboot has resolved more computer problems than there are grains of sand on a beach.

Creating a Backup Before making major changes to a system, always create a backup of the data. Hardware components can be replaced, but data often can't be. For that reason, always be vigilant about making data backups. Be sure to validate your backups by performing a test restore after you set them up. The only thing worse than not having a backup is thinking you have a backup only for it to fail.

Prioritizing Tasks Establish priorities—one user being unable to print to the printer of their choice isn't as important as a floor full of accountants unable to run payroll. Prioritize every job and escalate it (or de-escalate it) as you need to.

Documenting Your Process Last but perhaps most important, document everything—not just that there was a problem but also the solution that you found, the actions that you tried, and the outcome of each.

In the next section, we'll take you through general troubleshooting steps you can use as a guideline to follow.

Structured Troubleshooting Methodology

There's a famous quote attributed to Albert Einstein that states, "If I had an hour to solve a problem, I'd spend 55 minutes on the problem and 5 minutes on the solution." Whether or not he said that may be debatable, but the premise behind the quote is spot on because while this may seem obvious, it can't be overlooked: if you can't define the problem, you can't begin to solve it. This is true not only with computers but with every facet of life! Here are some good steps you can use to identify and resolve most problems you face:

Identify the problem. Sometimes, problems are relatively straightforward, but other times they're just a symptom of a bigger issue. For example, if a user isn't able to connect to the Internet from their computer, it could indeed be an issue with their system. But if other users are having similar problems, then the first user's difficulties might just be one example of the real problem.

Ask yourself, "Is there a problem?" Perhaps "the problem" is as simple as a customer expecting too much from the computer or not understanding how it's supposed to work.

The key to identifying the problem is to isolate the issue by figuring out what the user can and can't do. Tips to help with this include talking to the user ("Did anything change? How long has this been happening?") and gathering information. Again, isolate what's not working so you can focus on that.

Establish a theory. Now that you have gathered information, establish a theory of what is causing the problem. First question the obvious. Forgetting to check the obvious things can result in a long and unnecessary troubleshooting process. No video? It could be something to do with the monitor or the video card. Can't get to your favorite website? Is it that site? Is it your network card, the cable, your IP address, DNS server settings, or something else? Once you have defined the problem, establishing a theory about the cause of the problem—what is wrong—helps you develop possible solutions to the problem.

During this step, eliminate possibilities based on what is or isn't working. Also use external resources as needed, whether it be coworkers, Internet searches, manufacturers' websites, and the like.

Test the theory. You've eliminated possibilities and developed a theory as to what the problem is. Your theory may be specific, such as "the power cable is fried," or it may be a bit more general, like "we can't access the hard drive" or "there's a connectivity problem." No matter your theory, now is the time to start testing it. Again, if you're not sure where to begin to find a solution, the manufacturer's website is a good place to start!

Tips here include checking the simple stuff first, such as connections or rebooting the device. Also check to see if it may be a case of user error. See what the user is doing when the problem happens. Once you have tested the theory, you can devise the next steps to resolve the issue (if you haven't already done so).

Real World Scenario

"Is It Plugged In?" and Other Insulting Questions

Think about how you would feel if someone were to ask you this question. Your likely response would be, "Of course it is!" After all, you're not an idiot, right? You'll often get the same reaction to similar questions about the device being turned on. The problem is, making sure that it's plugged in and turned on are the first things that you should always do when investigating a problem.

When asking these types of questions, it's not what you say but how you say it. For example, instead of asking if it's plugged in, you could say something like, "Can you do me a favor and check to see that the cable is plugged in tightly? Maybe it got bumped." That generally gets the user at least to look at it without making them feel dumb. For power, something like, "What color are the lights on the front of the router? Are any of them blinking?" can work well.

Remember to ask neutral and nonthreatening questions. Make it sound like the computer is at fault, not the user. These types of things will help you build rapport and you'll be able to get more information so that you can solve problems faster.

Establish a plan of action. If your theory was right and the fix worked, then you're brilliant! If not, you need to look for the next option. After testing the theory, establish a plan of action to resolve the problem and implement the solution. This may take one of the following three paths:

- Verify full system functionality.
- If the first fix didn't work, try something else.
- If needed, implement the fix on other computers.

When trying solutions to fix a problem, make only one change to the computer at a time. If the change doesn't fix the problem, revert the system back to the way it was and then make your next change. Making more than one change at a time is not recommended for two reasons: One, you are never sure which change actually worked, and two, by making multiple changes at once, you might accidentally cause additional problems.

Verify functionality. After fixing the system, or all the systems affected by the problem, go back and verify full functionality. For example, if the users couldn't get to any network resources, check to make sure they can get to the Internet as well as to internal resources.

Some solutions may accidentally cause another problem on the system. For example, if you update software or drivers, you may inadvertently cause another application to have problems. There's obviously no way that you can or should test all applications on a computer after applying a fix, but know that these types of problems can occur. Just make sure that what you've fixed works and that there aren't any obvious signs of something else not working all of a sudden.

Another important thing to do at this time is to implement preventive measures, if possible. If it was a user error, ensure that the user understands ways to accomplish the task that won't cause the error to recur. If a cable melted because it was too close to a space heater under someone's desk, resolve the issue. If the computer overheated because an inch of dust was clogging the fan . . . you get the idea.

Document your findings. A lot of people can fix problems. But can you remember what you did when you fixed a problem a month ago? Maybe. Can one of your coworkers remember something you did to fix the same problem on that machine a month ago? Unlikely. Always document your work so that you or someone else can learn from the experience. Good documentation of past troubleshooting can save hours of stress in the future. Documentation can take a few different forms, but the two most common are personal and system-based.

We always recommend that technicians carry a personal notebook and take notes during the troubleshooting process. Some problems are long and complex, and it may be hard to remember which setting you changed 15 steps ago. The type of notebook doesn't matter—it can be digital or paper, big or small—use whatever works best for you. Take pictures with your phone if it will be helpful (if it's allowed). The notebook can be a lifesaver, especially when you're new to a job. Write down the problem, what you tried, and the solution. The next time you run across the same or a similar problem, you'll have a better idea of what to try. Eventually, you'll find yourself less and less reliant on your notebook, but it's incredibly handy to have.

System-based documentation is useful to both you and your coworkers. Many facilities have server logs of one type or another, conveniently located close to the machine. If someone makes a fix or a change, it gets noted in the log. If there's a problem, it's noted in the log

Troubleshooting Motherboards, RAM, CPUs, and Power Problems

To many who are not familiar with computers, that whirring, humming box sitting on or under their desk is an enigma. They know what shows up on the screen, where the power

button is, where to put a thumb drive, and what not to spill on their keyboard, but the insides are shrouded in mystery.

Fortunately for them, we're around. We can tell the difference between a hard drive and a motherboard and have a pretty good idea of what each part inside that box is supposed to do. When the computer doesn't work like it's supposed to, we can whip out our trusty screwdriver, crack the case, and perform surgery. And most of the time, we can get the system running just as good as new.

In the following sections, we're going to focus our troubleshooting efforts on the three key hardware components inside the case: motherboards, RAM, and CPUs, as well as power issues. These four components are absolutely critical to a computer system. Without a network card, you won't be able to surf the web. Without a processor, well, you won't be able to surf the web—or do much of anything else for that matter. So it makes sense to get started with motherboards, CPUs, RAM, and power. Problems with these separate components can often present similar symptoms, so it's good to discuss them all at the same time.

As you continue to learn and increase your troubleshooting experience, your value will increase as well. This is because, if nothing else, it will take you less time to accomplish common repairs. Your ability to troubleshoot from past experiences and gut feelings will make you more efficient and more valuable, which in turn will allow you to advance and earn a better income. We will give you some guidelines that you can use to evaluate common hardware issues that you're sure to face.

The following sections focus primarily on desktop computers. Mobile device and laptop-specific issues are covered in Volume 1 Chapter 13.

Identifying General Hardware Symptoms and Causes

Before we get into specific components, let's take a few minutes to talk about hardware symptoms and their causes at a general level. This discussion can apply to a lot of different hardware components.

Some hardware issues are pretty easy to identify. If flames are shooting out of the back of your computer, then it's probably the power supply. If the power light on your monitor doesn't turn on, it's the monitor itself, the power cord, or your power source. Other hardware symptoms are a bit more ambiguous. We'll now look at some general hardware-related symptoms and their possible causes.

Grinding and Other Unusual Noises

Have you ever been working on a computer and heard a noise that resembles fingernails on a chalkboard? If so, you will always remember that sound, along with the impending feeling of doom as the computer stops working.

Some noises on a computer are normal. The POST beep (which we'll talk about in a few pages) is a good sound. The whirring of a mechanical hard drive, case fans, video card fans, and power supply fan are familiar sounds. Some techs get so used to their "normal" system noises that if anything is slightly off pitch, they go digging for problems even if none are readily apparent.

A simple rule to remember about grinding and other random noises is this: for it to make a noise, it has to move. In other words, components with no moving parts (such as RAM, SSDs and CPUs) don't make sounds. Mechanical hard drives have motors that spin the platters. Power supply and case fans spin. Optical drives spin the discs. If you're hearing a grinding, whirring, scraping, or other unusual noise that you didn't expect, these are the likely culprits.

Exam Objective 5.1, "Given a scenario, troubleshoot motherboards, RAM, CPUs, and power," lists 12 symptoms you should know for the test. Unusual noise is one of them.

If you hear a whining sound and it seems to be fairly constant, it's more than likely a fan. Either it needs to be cleaned (desperately) or replaced. Power supplies that are failing can also sound louder and quieter intermittently because a fan will run at alternating speeds.

The "fingernails on a chalkboard" squealing could be an indicator that the heads in a mechanical hard drive have crashed into the platter. Thankfully, this isn't very common today, but it still happens. (Future generations of technicians will never know this sound, with the prevalence of SSDs today!) Note that this type of sound can also be caused by a power supply fan's motor binding up. A rhythmic ticking sound is also likely to be caused by a mechanical hard drive.

Problems with optical drives tend to be the easiest to diagnose. Those drives aren't constantly spinning unless you put some media in them. If you put a disc in and the drive makes a terrible noise, you have a good idea what's causing the problem.

So what do you do if you hear a terrible noise from the computer? If it's still responsive, shut it down normally as soon as possible. If it's not responsive, then shut off the power as quickly as you can. Examine the power supply to see if there are any obvious problems such as excessive dust, and clean it as needed. Power the system back on. If the noise was caused by the hard drive, odds are that the drive has failed and the system won't boot normally. You may need to replace some parts.

If the noise is mildly annoying but doesn't sound drastic, boot up the computer with the case off and listen. By getting up close and personal with the system, you can often tell where the noise is coming from and then troubleshoot or fix the appropriate part.

Never touch internal components when the case is off and the power is on! Doing so could result in a severe electrical shock to you and/or the components.

Overheating

Electronic components produce heat; it's a fact of life. While they're designed to withstand a certain amount of the heat that's produced, excessive heat can cause overheating and drastically shorten the life of components. There are two common ways to reduce heat-related problems in computers: heat sinks and cooling systems, such as case fans. Any component with its own processor will have a heat sink. Typically these look like big, finned hunks of

aluminum or another metal attached to the processor. Their job is to dissipate heat from the component so that it doesn't become too hot. Never run a processor without a heat sink! Nearly all video cards built today have GPUs with heat sinks as well.

Case fans are designed to take hot air from inside the case and blow it out of the case. There are many different designs, from simple motors to high-tech liquid-cooled models. Put your hand up to the back of your computer at the power supply fan and you should feel warm air. If there's nothing coming out, you either need to clean your fan out or replace your power supply. Some cases come with additional cooling fans to help dissipate heat. If your case has one, you should feel warm air coming from it as well.

Computers are like human beings; they have similar tolerances to heat and cold. In general, anything comfortable to us is comfortable to computers, although they do tend to like it colder than many of us do. They need a lot of clean, moving air in order to keep functioning.

Dust, dirt, grime, smoke, and other airborne particles can become caked on the inside of computers and cause overheating as well. This is most common in automotive and manufacturing environments. The contaminants create a film that coats the components, causing them to overheat and/or conduct electricity on their surface. Blowing out these exposed systems with a can of compressed air from time to time can prevent damage to the components. While you're cleaning the components, be sure to clean any cooling fans in the power supply or on the heat sink.

To clean the power supply fan, blow the air from the inside of the case. When you do this, the fan will blow the contaminants out the cooling vents. If you spray from the vents toward the inside of the box, you'll be blowing the dust and grime inside the case or back into the fan motor.

One way to ensure that dust and grime don't find their way into a desktop computer is to always leave the *blanks* (or slot covers) in the empty slots on the back of the case. Blanks are the pieces of metal or plastic that come with the case and cover the expansion slot openings. They are designed to keep dirt, dust, and other foreign matter from the inside of the computer. They also maintain proper airflow within the case to ensure that the computer doesn't overheat.

Real World Scenario

Creeping Chips

The inside of a computer is a harsh environment; the temperature inside the case of many computers is well over 100° F! When you turn on your computer, it heats up. Turn it off, and it cools down. After several hundred such cycles, some components can't handle the stress, and they begin to move out of their sockets. This phenomenon is known as *chip creep*, and it can be really frustrating.

Chip creep can affect any socketed device, including integrated circuits (ICs), RAM chips, and expansion cards. Many today have latches that hold them in place, but not all do. The solution to chip creep is simple: open the case and reseat the devices. It's surprising how often this is the solution to phantom problems of all sorts, particularly intermittent device failures, random reboots, and unexpected shutdowns.

Components that overheat a lot will have shorter lifespans. Sometimes they simply fail. Other times, they will cause *random shutdowns* before they fail. A PC that works for a few minutes and then locks up is probably experiencing *overheating* because of a heat sink or fan not functioning properly. To troubleshoot overheating, first check all fans inside the PC to ensure that they're operating, and make sure that any heat sinks are firmly attached to their chips.

Overheating and random shutdowns are symptoms specified in Exam Objective 5.1.

In a properly designed, properly assembled PC case, air flows in a specific path driven by the power supply fan and using the power supply's vent holes. Make sure that you know the direction of flow and that there are limited obstructions and no dust buildup. Cases are also designed to cool by making the air flow in a certain way. Therefore, operating a PC with the cover removed can make a PC more susceptible to overheating, even though it's "getting more air."

Although CPUs are the most common component to overheat, video cards are also quite susceptible, especially for those with high-end graphics needs such as video producers, graphics designers, and gamers. Occasionally other chips on the motherboard—such as the RAM or chips on other devices—may also overheat. Extra heat sinks, fans, or higher-end cooling systems may be installed to cool these chips.

If the system is using a liquid cooling system, know that they have their own set of issues. The pump that moves the liquid through the tubing and heat sinks can become obstructed or simply fail. If this happens, the liquid's temperature will eventually equalize with that of the CPU and other components, resulting in their damage. Dust in the heat sinks has the same effect as with non-liquid cooling systems, so keep these components as clean as you would any such components. Check regularly for signs of leaks that might be starting and try to catch them before they result in damage to the system.

Burning Smells, Smoke, and Visible Damage

A *burning smell* or smoke coming from your computer is never a good thing. While it normally gets warm inside a computer case, it should never be hot enough inside there to melt plastic components or cause visible damage, but it does happen from time to time. And power problems can sometimes cause components to get hot enough to smoke.

A burning smell is a specific symptom called out in Exam Objective 5.1.

If you smell an odd odor or see smoke coming from a computer, shut it down immediately. Open the case and start looking for visible signs of damage. Things to look for include melted plastic components and burn marks on circuit boards. The good news about visible damage is that you can usually figure out which component is damaged pretty quickly. The bad news is that it often means you need to replace parts.

Visible damage to the outside of the case or the monitor casing might not matter much as long as the device still works. But if you're looking inside a case and see burn marks or melted components, that's a sure sign of a problem. Replace damaged circuit boards or melted plastic components immediately. After replacing the part, it's a good idea to monitor the new component for a while too. The power supply could be causing the problem. If the new part fries quickly too, it's time to replace the power supply as well.

Looking for Lights and Messages

Many hardware devices have status indicator lights that can help you identify operational features or problems with a device. Obviously, when you power on a system, you expect the power light to come on. If it doesn't, you have a problem. The same holds true for other external devices, such as wireless routers, external hard drives, and printers. In situations in which the power light doesn't come on and the device has no power, always obey the first rule of troubleshooting: Check your connections first!

Beyond power indicators, several types of devices have additional lights that can help you troubleshoot. If you have a switch, hub, or other connectivity device, you should have an indicator for each port that lights up when there is a connection. Some devices will give you a green light for a good connection and a yellow or red light if they detect a problem. A lot of connectivity devices will also have an indicator that blinks or flashes when traffic is going through the port. Sometimes it's the same light that indicates a connection, but at other times it's a separate indicator. The same holds true for NICs. They usually have a connectivity light and a transmission light. If no lights are illuminated, it can indicate a lack of connection.

Many computers also have hard drive activity lights. When disk reads or writes occur, the light will blink; otherwise, it will be off. A hard drive indicator that is constantly on is generally not a good sign; it could indicate that the hard drive is constantly busy or that the system is frozen, either of which is bad.

Most keyboards will have status lights for the Caps Lock and Num Lock keys. If you believe a system is locked up, try pressing the Caps Lock or Num Lock key on the keyboard to see if the lights change. If they don't, that's a sign that the system is unresponsive. If no lights ever illuminate on the keyboard, it could be that the keyboard is disconnected or that there is a system power issue.

In addition to lights, you should look for alerts and error messages when possible. An alert is a message generated by a hardware device. In some cases, the device has a display

panel that will tell you what the alert is. A good example of this is an office printer. Many have an LCD display that can tell you if something is wrong.

Other alerts and error messages will pop up on the computer screen. If the device is attached to a specific computer, the alert will generally pop up on that computer's screen. Some devices can be configured to send an alert to a specific user account or system administrator, so the administrator will get the alert regardless of which computer they are currently logged into.

Operating system error messages and alerts are generally logged in an event log, such as Windows Event Viewer. Troubleshooting operating system errors is an objective on the 220-1202 exam, and we cover Event Viewer in Volume 2 Chapter 2, "Windows Configuration."

If you have a device with lights and you're not sure what they mean, or if you're receiving an error message that is confusing, it's best to check the manual or the manufacturer's website to learn about them. The manufacturer's website is generally a great place to go for troubleshooting tips!

Intermittent Device Failure

Intermittent problems are absolutely the worst to deal with. They are frustrating for technicians because the system will inevitably work properly when the tech is there to fix it. The users also get frustrated because they see the problem happen, but, of course, it works fine when the tech shows up!

Treat intermittent failures just as you would a persistent issue, if at all possible. See if there were any error messages, or if it happens when the user tries a certain action. Maybe it occurs only when the system has been on for a while or when a specific application is open. Try to narrow it down as much as possible. In many cases, an intermittent failure means that the device is slowly but surely dying and needs to be replaced. If it's something obvious, such as a network card or a disk read/write failure, you know what to replace. If not, and it's something random like intermittent lockups, it may take trial-and-error to find the right part to replace, especially if there were no error messages. Intermittent or unexpected lockups or shutdowns may be a motherboard, CPU, or RAM problem. If you have nothing else to go on, try replacing one at a time to see if that resolves the issue.

Identifying BIOS/UEFI and POST Routine Problems

Every computer has a Basic Input/Output System (BIOS) or Unified Extensible Firmware Interface (UEFI) that acts as an interface between the computer's hardware and any operating system installed on that hardware. It's the first software (technically, firmware) that loads as a system boots, and as such, it plays a critical role in computer operation. If the BIOS/UEFI doesn't work properly, the user will never get to the point where their operating system starts. In this section, we will cover issues specific to the BIOS/UEFI, as well as the boot process it controls, which is the *power-on self-test (POST)*.

BIOS/UEFI Issues

As already mentioned, the BIOS/UEFI is the layer between the hardware and the operating system. Older systems may have a BIOS; UEFI is newer and has more features.

Out-of-Date BIOS/UEFI First, computer BIOS/UEFIs don't go bad; they just become out-of-date. This isn't necessarily a critical issue; they will continue to support the hardware that came with the box. It does, however, become an issue when the BIOS doesn't support some component that you would like to install—virtualization, for instance.

Most of today's BIOS/UEFIs are written to an electrically erasable programmable read-only memory (EEPROM) and can be updated through the use of software. This process is called *flashing the BIOS*. Each manufacturer has its own method for accomplishing this. Check the documentation for complete details.

If you make a mistake in flashing the BIOS/UEFI, the computer can become unbootable. If this happens, your only option may be to ship the box to a manufacturer-approved service center. Be careful!

Losing a Computer's Settings A fairly common issue with the BIOS/UEFI is when it fails to retain your computer's settings, such as date/time and hard drive settings. The BIOS/UEFI uses a small battery (much like a watch battery), called the *CMOS battery*, on the motherboard to help it retain settings when the system power is off. If this battery fails, the BIOS/UEFI won't retain its settings. Simply replace the battery to solve the problem.

Inaccurate system date/time is a specific symptom called out in the Exam Objective 5.1.

Checking the Boot Priority Finally, remember that your BIOS/UEFI also contains the boot priority (also sometimes called the boot sequence) for your system. You probably boot to the first hard drive in your system (the one that contains the OS boot files), but you can also set your BIOS/UEFI to boot from a secondary hard drive, an optical drive, a USB port, or the network. If your computer can't find a proper boot device, it could be that it's attempting to boot to an incorrect device. Check the BIOS/UEFI to see if you need to change the boot sequence. To do this, perform the following steps:

1. Reboot the system and look for the message telling you to press a certain key to enter the BIOS/UEFI (usually something like F2).
2. Once you're in the BIOS/UEFI, find the menu with the boot priority (like the one shown in Figure 12.1) and set it to the desired order.

If the changes don't hold the next time you reboot, check the battery.

FIGURE 12.1 UEFI boot priority settings

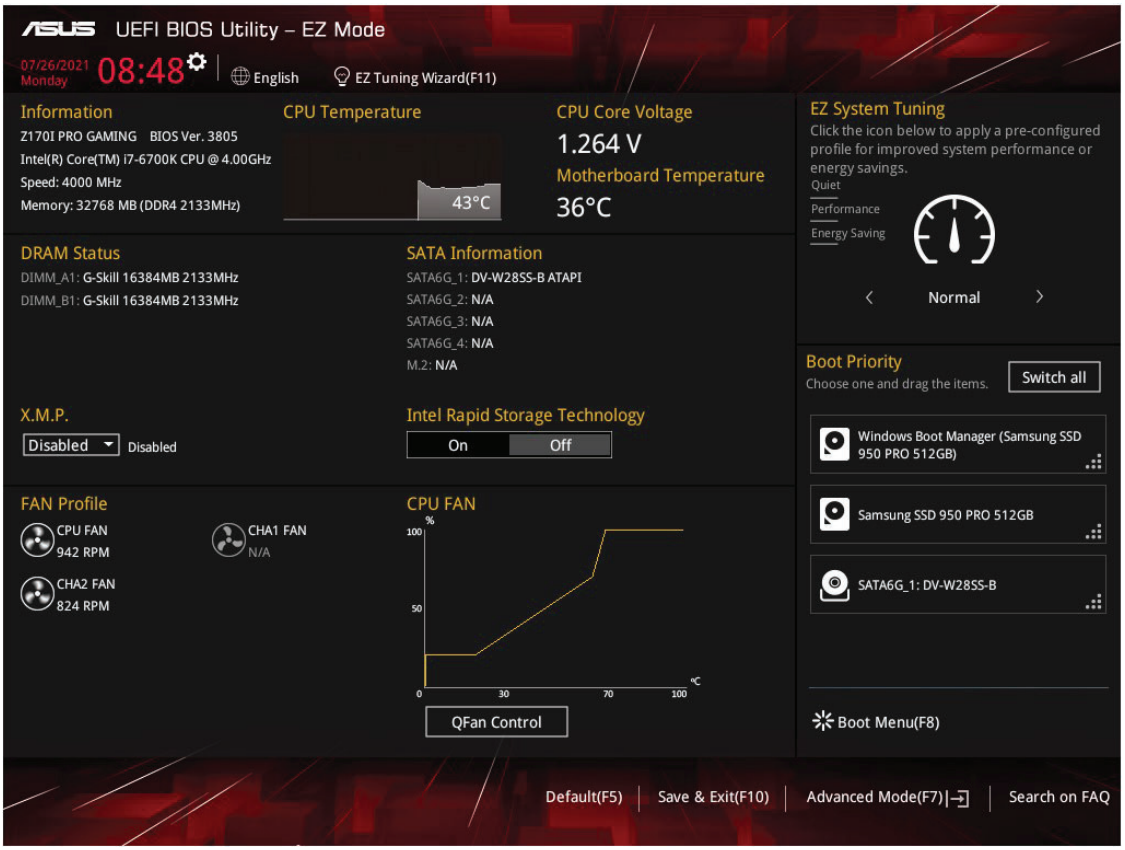

There may be times that you want to boot to an alternate device, such as a USB drive. During the part of the boot process when the BIOS/UEFI entry key is shown, most systems will also show you an alternate key to press (such as the spacebar) to boot a different device. This can be very helpful when you're troubleshooting a system that won't load the OS properly, as long as you have the installation medium handy.

POST Routine Problems

Every computer has a diagnostic program built into its BIOS/UEFI called the POST. When you turn on the computer, it executes this set of diagnostics. Many steps are involved in the POST, but they happen very quickly, they're invisible to the user, and they vary among BIOS/UEFI vendors. The steps include checking the CPU, checking the RAM, checking for the presence of a video card, and verifying basic hardware functionality. The main reason to be aware of the POST's existence is that if it encounters a problem, the boot process stops. Being able to determine at what point the problem occurred could help you troubleshoot.

If the computer doesn't perform the POST as it should, one way to determine the source of a problem is to listen for POST code beeps, also known as a *beep code*. This is a series of beeps from the computer's speaker. A successful POST generally produces a single beep. If there's more than one beep, the number, duration, and pattern of the beeps can sometimes tell you what component is causing the problem. However, the beeps differ depending on the BIOS manufacturer and version, so you must look up the beep code in a chart for your particular BIOS. AMI BIOS, for example, relies on the number of beeps and uses patterns of short and long beeps. Unfortunately, not all computers today give any beep codes because they don't contain the internal piezoelectric speaker.

Power-on self-test (POST) beeps is a specific symptom called out in Exam Objective 5.1.

Another way to determine a problem during the POST routine is to use a *POST card*. This is a circuit board that fits into an expansion slot (USB, PCIe, or PCI) in the system and reports numeric codes as the boot process progresses. Each code corresponds to a particular component being checked. If the POST card stops at a certain number, you can look up that number in the manual for the card to determine the problem. USB POST cards, like the one shown in Figure 12.2, are easy to use—you don't have to crack the case to check for POST errors—and they can be used to test laptops as well.

FIGURE 12.2 USB POST card

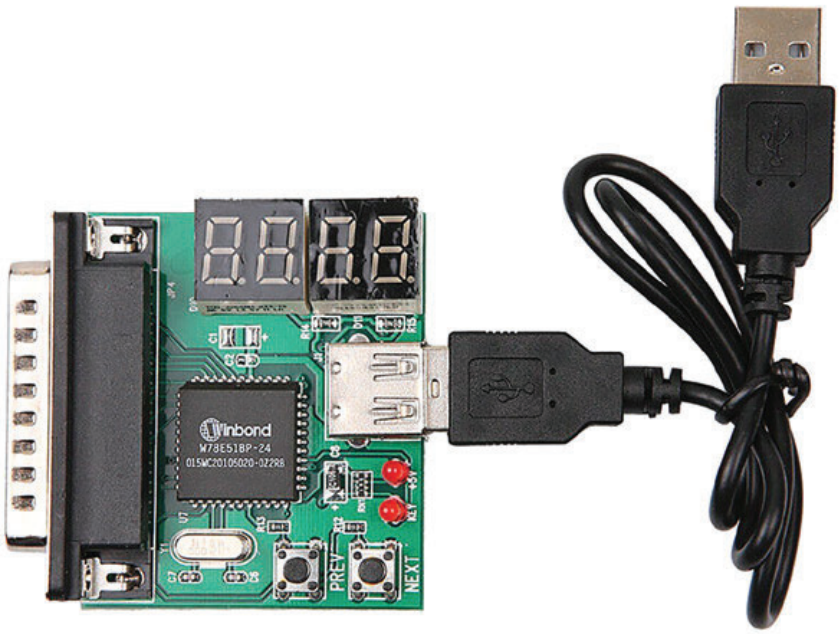

Identifying Motherboard and CPU Problems

Most motherboard and CPU problems manifest themselves by the system appearing to be completely dead. However, "completely dead" can be a symptom of a wide variety of problems, not only with the CPU or motherboard but also with the RAM or the power supply. At other times, a failing motherboard or CPU will cause the system to lock up completely, or "hang," requiring a hard reboot, or the failing motherboard or CPU may cause continuous reboots. A POST card may be helpful in narrowing down the exact component that is faulty.

When a motherboard fails, it's usually because it has been damaged. Most technicians can't repair motherboard damage; the motherboard must be replaced. Motherboards can become damaged due to physical trauma, exposure to electrostatic discharge (ESD), or short-circuiting. To minimize the risk, observe the following rules:

- Handle a motherboard as little as possible, and keep it in an antistatic bag whenever it's removed from the PC case.
- Keep all liquids well away from the motherboard; water can cause a short circuit.
- Wear an antistatic wrist strap when handling or touching a motherboard.
- When installing a motherboard in a case, make sure that you use brass standoffs with paper or plastic washers to prevent any stray solder around the screw holes from causing a short circuit with the metal of the screw.

A CPU may fail because of physical trauma or short-circuiting, but the most common cause for a CPU not to work is overheating, and most overheating issues are due to installation failures. This means that the heat sink and/or fan must be installed properly along with the processor. For example, air gaps in the thermal paste between the CPU and heat sink can cause the processor to run hot and eventually burn up. With a PGA- or LGA-style CPU, ensure that the CPU is oriented correctly in the socket. With an older SECC- or ZIF-style CPU, make sure the CPU is completely inserted into its slot or socket.

I/O Port and Cable Problems

Input/output (I/O) ports are most often built into the motherboard and include USB as well as legacy parallel and serial ports. All of them are used to connect external peripherals to the motherboard. When a port doesn't appear to be functioning, make sure the following conditions are met:

- The cables are snugly connected.
- The port has not been disabled in BIOS/UEFI Setup.
- The port has not been disabled in Device Manager in Windows.
- No pins are broken or bent on the male end of the port or on the cable being plugged into it.

If you suspect that it's the port, you can purchase a loopback plug to test its functionality. If you suspect that the cable, rather than the port, may be the problem, swap out the cable with a known good one. If you don't have an extra cable, you can test the existing cable with a multimeter by setting it to ohms and checking the resistance between one end of the cable and the other.

Use a pin-out diagram, if available, to determine which pin matches up to which at the other end. There is often—but not always—an inverse relationship between the ends. In other words, at one end pin 1 is at the left, and at the other end it's at the right on the same row of pins. You see this characteristic with D-sub connectors where one end of the cable is male and the other end is female.

Blank Screen

Sometimes you will run into a computer with no video output—a *blank screen*. This symptom could be the fault of a few different components. Recall that for video to be produced, there needs to be a video card, cable, and a screen, and the screen may have several components such as the screen itself, inverter, backlight, and others. We'll cover it here, though, because with so many motherboards today having video circuitry built into them, it could very well be the motherboard.

Blank screen is a specific symptom called out in Exam Objective 5.1.

If there's no video, as always, check the obvious first. Is the monitor plugged in and turned on? Does it appear to be getting a signal from the video card? Does the light on the monitor make it look like the monitor has gone into sleep mode? If so, perhaps turning the monitor off and back on will do the trick.

If you've checked the usual suspects, try a different monitor and video cable. If there's still no video, then you may try a different video card. On almost all motherboards with built-in video circuitry, the onboard video electronics will be disabled when an expansion video card is installed. Of course, make sure you're plugging the monitor into the new video card and not the video port on the motherboard! (It's happened to all of us, and yes, it can be kind of embarrassing.) If the video circuitry on the motherboard is faulty, either use an expansion card or replace the board.

Capacitor Swelling

Many motherboards have capacitors on them, which store electricity. They are short cylindrical tubes. Sometimes, when capacitors fail, they will swell and brownish-red electrolyte residue may seep out of the vents in the top—an example is shown in Figure 12.3. These are called *distended capacitors*, also known as capacitor swelling.

Capacitor swelling is a specific symptom called out in Exam Objective 5.1.

If a capacitor fails, the motherboard will not work. You have a couple of options:

- The first and probably best option is to replace the motherboard. Whatever you do, do *not* touch the residue coming from a distended capacitor! It can cause serious chemical burns.
- The second option is to drain the energy from the failed capacitor and replace it. Do this only if you have specialized training on how to deal safely with capacitors, because they can cause lethal shocks.

FIGURE 12.3 Distended capacitors on a motherboard

Source: Bushytails / Wikimedia Commons / CC BY SA 3.0

Identifying Memory Issues

Isolating memory issues on a computer is one of the most difficult tasks to do properly because so many memory problems manifest themselves as software issues. For example, memory problems can cause application crashes and produce error messages such as general protection faults (GPFs). Memory issues can also cause a fatal error in your operating system, producing proprietary crash screens such as the infamous *Blue Screen of Death (BSOD)* in Windows or the rotating *pinwheel* in macOS. Sometimes these are caused by the physical memory failing. At other times, they are caused by bad programming, when an application writes into a memory space reserved for the operating system or another application.

Application crashes and proprietary crash screens are specific symptoms mentioned in A+ Exam Objective 5.1.

In short, physical memory problems can cause app and system lockups, unexpected shutdowns or reboots, or the errors mentioned in the preceding paragraph. They can be challenging to pin down. If you do get an error message related to memory, be sure to write down the memory address if the error message gives you one. If the error happens again, write down the memory address again. If it's the same or a similar address, then it's very possible that the physical memory is failing. You can also use one of several hardware- or software-based RAM testers to see if your memory is working properly. Sometimes switching the slot that the RAM is in will help, but more often than not the RAM needs to be replaced.

Memory issues can also be caused by the virtual memory, which is an area of the hard drive set aside to emulate memory. The operating system creates and manages a paging file (in Windows, it's called `PAGEFILE.SYS`) on the hard drive to act as memory when the system needs more than what the physical RAM can provide; oftentimes, this paging file is dynamic in size. If the hard drive runs out of room for the paging file, memory issues can appear, or the system may have sluggish performance. As a rule of thumb, ensure that at least 10 percent of the hard drive space is free.

This Computer Is SO Slow . . .

Sluggish performance is a specific symptom mentioned in A+ Exam Objective 5.1. It's a tricky problem to isolate because it can be caused by many different things. For example, it could be one of the following:

- Low physical or virtual memory
- Low available space on the primary hard drive
- Failing motherboard, CPU, or hard drive
- Poorly coded software apps
- Too many apps open at once (causing low memory)
- Malware on the computer

The easy first step to take when troubleshooting a sluggish Windows-based computer is to look at memory usage in Task Manager. There are several ways to open Task Manager, such as pressing Ctrl+Alt+Delete and then choosing Task Manager from the list, right-clicking the Taskbar and selecting Task Manager, or pressing Ctrl+Shift+Esc. Task Manager's Processes page is shown in Figure 12.4. This gives real-time statistics on how busy the CPU, RAM, hard drive, and network connection are. If the CPU, memory, or disk counters are over 80 percent, the system may appear sluggish. Clicking the column header (where the total percent and counter name is) will sort it, so you can see which app or process is causing the most grief. If there are a lot of open apps, closing them may help. If the system isn't running many apps, it may be a background process or a Windows process. Before closing a background or Windows process down, do an Internet search to see what it does—randomly closing processes can limit system functionality (until the computer is rebooted) or cause a system crash.

FIGURE 12.4 Windows Task Manager

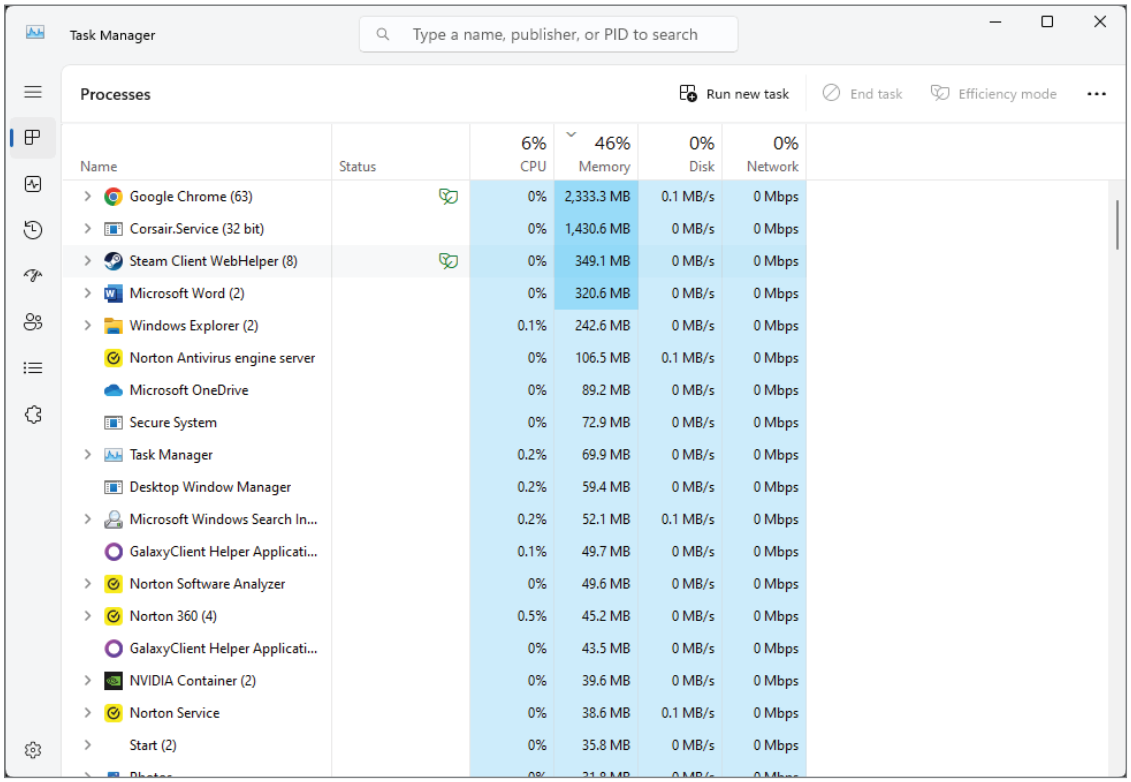
Task Manager
Type a name, publisher, or PID to search
Processes
Run new task
End task
Efficiency mode
Name
Status
6% CPU
46% Memory
0% Disk
0% Network
Google Chrome (63) 0% 2,333.3 MB 0.1 MB/s 0 Mbps
Corsair.Service (32 bit) 0% 1,430.6 MB 0 MB/s 0 Mbps
Steam Client WebHelper (8) 0% 349.1 MB 0 MB/s 0 Mbps
Microsoft Word (2) 0% 320.6 MB 0 MB/s 0 Mbps
Windows Explorer (2) 0.1% 242.6 MB 0 MB/s 0 Mbps
Norton Antivirus engine server 0% 106.5 MB 0.1 MB/s 0 Mbps
Microsoft OneDrive 0% 89.2 MB 0 MB/s 0 Mbps
Secure System 0% 72.9 MB 0 MB/s 0 Mbps
Task Manager 0.2% 69.9 MB 0 MB/s 0 Mbps
Desktop Window Manager 0.2% 59.4 MB 0 MB/s 0 Mbps
Microsoft Windows Search In... 0.2% 52.1 MB 0.1 MB/s 0 Mbps
GalaxyClient Helper Applicati... 0.1% 49.7 MB 0 MB/s 0 Mbps
Norton Software Analyzer 0% 49.6 MB 0 MB/s 0 Mbps
Norton 360 (4) 0.5% 45.2 MB 0 MB/s 0 Mbps
GalaxyClient Helper Applicati... 0% 43.5 MB 0 MB/s 0 Mbps
NVIDIA Container (2) 0% 39.6 MB 0 MB/s 0 Mbps
Norton Service 0% 38.6 MB 0.1 MB/s 0 Mbps
Start (2) 0% 35.8 MB 0 MB/s 0 Mbps

A second thing to check is free disk space. Once the primary drive gets to be under 10 percent free, the system can become slow to respond. Check this by looking at the hard drive properties (open Windows Explorer, right-click the hard drive, then choose Properties), as shown in Figure 12.5. If the hard drive is too full, you can use the Disk Cleanup utility to free space. Or, you can manually delete unneeded files or install a second drive, and copy apps, photos, videos, and other items to the secondary drive.

FIGURE 12.5 Local disk properties

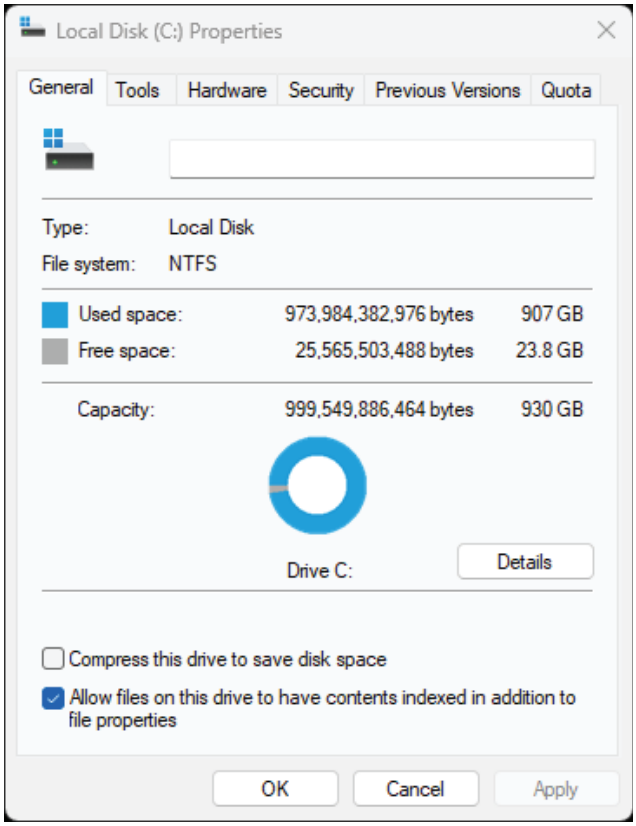

Identifying Power Supply Problems

Power supply problems can manifest themselves in two ways. In the first, you will see an obvious problem, such as an electrical flash or possibly a fire. In the second, the system doesn't respond in any way when the power is turned on. Hopefully you don't have to deal with many of the first type!

Recall from Volume 1 Chapter 2, "Expansion Cards, Storage Devices, and Power Supplies," that some power supplies have a voltage switch on the back to change from 110v–120v power to 220v–240v power to meet different countries' power specifications. If a computer has been moved from one country to another, make sure that it's on the right setting before you plug it in and try to power it up.

When the system doesn't respond ("no power") when you try to power it up, make sure the outlet is functional and try a new power cable. If those check out, open the case, remove the power supply, and replace it with a new one. Partial failures, or intermittent power

supply problems, are much less simple. A completely failed power supply gives the same symptoms as a malfunctioning wall socket, uninterruptible power supply (UPS), or power strip; a power cord that is not securely seated; or some motherboard shorts (such as those caused by an improperly seated expansion card, memory stick, CPU, and the like). You want to rule out those items before you replace the power supply and find that you still have the same problem as when you started.

No power is a specific symptom mentioned in A+ Exam Objective 5.1.

At other times, the power supply fan might spin but the rest of the system does not appear to get power. This can be a power supply issue or possibly a motherboard issue. (Recall that the power supply plugs into the motherboard, and several devices don't have a power cable but power themselves from the motherboard.) Be aware that different cases have different types of on/off switches. The process of replacing a power supply is a lot easier if you purchase a replacement with the same mechanism.

Real World Scenario

Hot, Hot, Hot

Several years ago, the company with which one of the authors was working got in a batch of hardware that it had purchased from another company. He and another tech were building Frankensteins out of the plethora of parts they had.

They put RAM into one of the systems and powered it on. Immediately there was an electrical arc from the RAM to the motherboard, so they shut it back off. The arc was present for a split second, and they had the box powered down within a second or two after that.

The RAM module had a pretty obvious burn mark on it, so the author went to take it out and promptly scorched his fingers when he touched it. It was searing hot! They let it cool down for about 20 minutes before going back to take it out. The moral of the story: be careful not to burn yourself on fried components.

Incidentally, they put a new motherboard and new RAM into the same case and powered it up only to see the exact same thing happen. Fried. (Fortunately, the author was smart enough not to burn himself a second time!) The verdict? Bad power supply. After replacing the power supply and trying a third motherboard and RAM combination, they had a functioning system.

If you're curious as to the state of your power supply, you can buy hardware-based power supply testers online starting at about $10 and running up to several hundred dollars. Multimeters are also effective devices for testing your power supplies.

Never try to repair or disassemble a power supply. They contain capacitors that can store a lethal amount of electricity, even when they are powered off and have been unplugged. The high risk of electrocution and the relatively low cost of a new power supply make working on them something to avoid.

Exercise 12.1 walks you through the steps of troubleshooting a few specific hardware problems. The exercise will probably end up being a mental one for you, unless you have the exact problem that we're describing here. As practice, you can write down the steps that you would take to solve the problem and then check to see how close you came to our steps. Clearly, there are several ways to approach a problem, so you might use a slightly different process, but the general approach should be similar. Finally, when you have found the problem, you can stop. As you go through each step, assume that it didn't solve the issue so you need to move on to the next step.

For additional troubleshooting experience, you can watch videos on YouTube or `https://iFixIt.com`. There are many great examples for hundreds of different types of problems.

EXERCISE 12.1

Troubleshooting Practice

Issue 1: Blank screen on bootup. You turn the computer on, and there's nothing on the screen.

1. Check to make sure the monitor is on. Is its power light on?

   Seriously. Check it. Sometimes 5 seconds of checking the obvious can save you an hour of wasted time.

2. Is the monitor getting a signal?

   Some monitors will go into sleep mode if they don't get a signal. Check the connections. If all the connections are good and you're not getting a signal, it could be the video card.

3. Did the system POST properly? Did you get a POST beep or a beep code?

   No POST likely indicates a bigger problem than just the video card or the monitor. If you do get a POST beep but never see anything, try a different monitor.

EXERCISE 12.1 *(continued)*

4. Did you ever see anything on the screen? BIOS/UEFI information? Did the OS start to load and then go blank?

   The key to troubleshooting an ambiguous situation like this is to ask, "What is the last thing that worked as it was supposed to?" That will help you determine what you need to fix.

Issue 2: The power supply fan spins, but no other devices have power.

1. Did you hear a POST beep or a beep code?

   Odds are you that didn't get any sounds, but it's always good to reboot and double-check.

2. Disconnect all internal and external peripherals so that the only component drawing power is the motherboard (with CPU and RAM, of course). Does it POST then?

   If you disconnect everything and it still doesn't POST, odds are that your motherboard is fried. If it POSTs, then start plugging components back in one at a time, starting with your hard drive and other internal devices and working your way to the external peripherals. You'll eventually get to the part that's causing the problem.

3. If you have a power supply tester or multimeter, now would be a good time to make sure that the power supply is working properly. There's no sense in replacing components, such as the motherboard, if the power supply is just going to fry them!

   Again, with all troubleshooting, it's imperative to narrow down the problem to isolate the cause. If you can do that, then fixing it should be the easy part.

Troubleshooting Drives and RAID Arrays

Even though storage drives aren't needed, persistent storage is almost universally included in computing devices today. And when they don't work, users tend to get upset. Losing a hard drive's worth of data is incredibly frustrating, especially if there's no suitable backup.

Storage drives present unique problems simply due to their nature—there are multiple technologies in use today. Some of them are devices with moving parts, which means that they are more prone to mechanical failure than a motherboard or a stick of RAM. Others function essentially like RAM and plug directly into a socket instead of needing a cable. If they fail, there's not much you can do besides replace them.

Before we get into common symptoms and solutions, remember that storage system problems usually stem from one of the following three causes:

- The adapter (that is, the M.2, SATA, or other interface) is bad.
- The hard drive (or optical disc) is bad or failing.
- The adapter and disk are connected incorrectly.

The first and last causes are easy to identify, because in either case, the symptom will be obvious: the drive won't work. You won't be able to get the computer to communicate with the disk drive. The way to see which component is at fault is to disconnect and reconnect, or to try the device in another system (or try another drive in the affected system). However, if the problem is a bad or failing disk drive, the symptoms aren't always as obvious. Those are problems we will need to dive deeper into. In the following sections, we'll discuss hard disk problems, including using S.M.A.R.T. technology and RAID arrays. Then we'll finish by taking a quick look at optical drive issues.

Lights and Sounds

While looking at *light-emitting diode (LED) status indicators* is listed as a subobjective specifically under storage drives, it could be put in several different places as well. A lot of devices have lights that can indicate whether or not a component is working. Storage systems usually have some sort of activity indicator that blinks when the drive is busy either reading or writing data. If the light never comes on, or if the light is on constantly without flickering, there could be a problem.

External network attached storage (NAS) and redundant array of independent (or inexpensive) disks (RAID) storage enclosures have lights as well and may have many more than a standard desktop or laptop computer. For example, many stand-alone RAID arrays have a light that only illuminates if a drive has failed and needs to be replaced. We'll get into RAID in more detail later in this section. The point is, look for indicator lights and understand what they're communicating to you.

Some external NAS and RAID devices will also have audible alarms, such as a siren or a series of beeps, that sound if a drive fails or other conditions arise that need the attention of the administrator. Consult the device manufacturer's website or service manual to understand what the alarm is telling you.

Exam Objective 5.2, "Given a scenario, troubleshoot drive and RAID issues" lists 12 common symptoms to understand. They include these four:

- Light-emitting diode (LED) status indicators
- Grinding noises
- Clicking noises
- Audible alarms

Storage devices that have moving components will make sounds. Mechanical HDDs have a whirring sound as the platters spin, and when reading and writing they make an irregular clicking, ticking, or scratching sound. An optical drive spins up when a disc is inserted, and it too will whirr. SSDs have no moving parts, and therefore they do not make sounds. (Well, technically, if one were to get fried it could make an electrical pop. If you hear that, the drive is toast.)

A grinding noise from a mechanical storage device is a very bad sign. That means there is a failure in the motor or spindle, or if it sounds more like fingernails on a chalkboard, it means the read/write heads have crashed into the platter and are cutting grooves into it. You'll only need to hear that sound once to remember it forever. If the drive is still operational, get all important data off the drive immediately and replace the drive. A regular, rhythmic ticking or clicking sound is bad too—that usually means the drive is failing or has failed. The solution is the same as for a grinding noise. If you can, get important data off the drive and then replace it.

Talking about hard drive failures is a good time to remind everyone about the importance of regular backups. If the data is irreplaceable, then it needs to be backed up. More frequently updated data needs to be updated more often, for the backups to be current and helpful. On the flip side, after a drive fails is probably not the best time to lecture someone on not having backed up their data. Many years ago, one of the authors took a service call from a third-year PhD student who had a bad hard drive. The student's dissertation—the only copy of it—was on the drive. A bit of empathy was in order, considering the circumstances.

Devices Not Found

If a storage device is plugged in and working, the BIOS/UEFI should detect it first, and then as the operating system loads, it will recognize the drive as well. If someone gives you the symptom that their hard drive isn't being found, the first thing to do is clarify where it's not being found. Is it the BIOS/UEFI, or in the operating system?

Bootable device not found and missing drives in OS are both sub-objectives of Exam Objective 5.2. Be sure to understand how to troubleshoot each scenario.

Bootable Device Not Found This could manifest itself in a few different ways, such as a complete failure to boot, the hard drive not being recognized by the BIOS/UEFI, or the OS not being found. Failure to boot at all likely means the drive is not properly connected or it's dead. Do your due diligence and reseat your connections and try different cables, or try the drive in another machine if possible. Most BIOSs/UEFIs today autodetect the hard drive. If that auto-detection fails, it's bad news for the hard drive, unless there's a cable or connection issue.

Finally, a system that boots fine but can't find the OS could indicate a problem with the master boot record (MBR) or boot sector on the hard drive. To fix this in any current version of Windows, boot to bootable media (USB or optical disc) and enter the Windows Recovery Environment (WinRE). In WinRE, you can get to a command

prompt and use `bootrec /fixmbr` to fix the MBR and `bootrec /fixboot` to fix the boot sector.

Failed reads and writes from hard drives can also cause the operating system to crash, resulting in a proprietary crash screen (BSOD in Windows or pinwheel in macOS) error.

Missing Drives in OS Sometimes you have a storage device installed, but you are unable to find it in Windows File Explorer or the Finder in macOS. This could be a connection issue, or perhaps the drive hasn't been properly set up for use. Windows comes with a utility called Disk Management, and Macs have Disk Utility to perform storage device management. If those utilities don't see the drive or mark it as unusable, there are a few things to try:

- Check the connections and try a new cable or port.
- Double-check that the drive is found by the BIOS/UEFI and enabled. If the BIOS/UEFI doesn't see it, the OS won't either.
- Ensure that the drive is found in Disk Management or Disk Utility. If the BIOS/UEFI sees the drive, then the disk utilities should too—but that doesn't necessarily mean the drive is ready for use by File Explorer or the Finder.
- Prepare the drive by initializing it (if needed), allocate space by creating a partition, and assign a drive letter. Without a partition and drive letter, File Explorer and the Finder won't be able to use the space.
- Update the storage device's driver, if possible. This is admittedly a long shot, but it could work, particularly if you're dealing with older hardware.
- Check the drive for bad sectors or other damage. If enough sectors are damaged, the operating system may deem the drive unreadable.

Disk Management is covered in Volume 2 Chapter 2, and other software utilities used to manage hard drives such as `chkdsk`, `format`, and `diskpart` are covered in Volume 2 Chapter 3, "Windows Administration."

Performance Issues

In this category are times when the drive is working, but perhaps not as well as it should or once did. The A+ exam objectives that fall into this group include the following:

- Data loss/corruption
- Extended read/write times
- Low performance input/output operations per second (IOPS)

The three performance issues listed, data loss/corruption, extended read/write times, and low performance input/output operations per second (IOPS), are explicitly mentioned in Exam Objective 5.2.

A failing hard drive might exhibit data loss or corruption or very slow (extended) read/write times. They can also be a symptom of the hard drive being too full. Hard drives move information around a lot, especially temporary files. If the drive doesn't have enough free space (at least 10 percent), it can slow down dramatically. The solution here is to remove files or old applications to free up space and look at defragmenting the hard drive. If problems persist, consider formatting the hard drive and reinstalling the OS. If the issues don't go away, assume that the hard drive is on its last legs. Exercise 12.2 shows you a few different ways to look at storage drives, freeing up space, and optimizing drives.

EXERCISE 12.2

Viewing and Managing Storage Space

Windows 11 has built-in utilities to view and manage available storage space. In this exercise, you'll learn how to look at available space and optimize or clean up drives.

1. Open the Storage app by choosing Start ➢ Settings ➢ System ➢ Storage. You will see a screen similar to the one shown in Figure 12.6. (You can also get to an adjacent place through the drive's Properties page that you saw in Figure 12.5 by clicking the Details button.)

FIGURE 12.6 Storage app

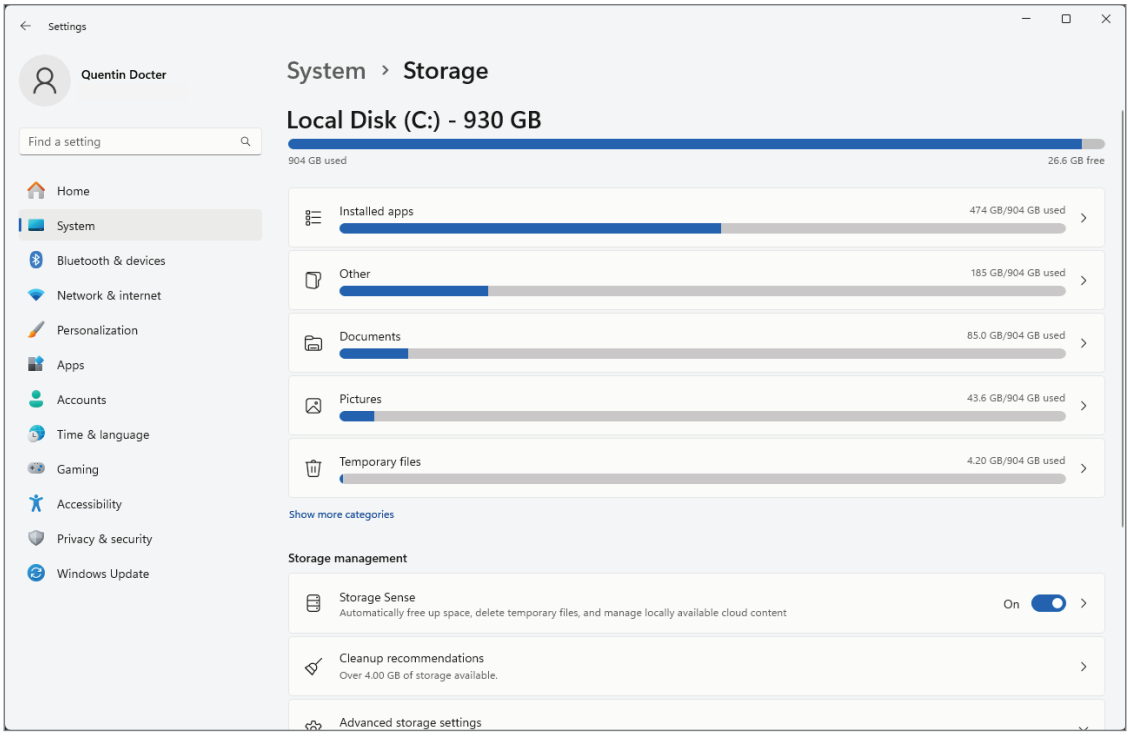

Here you will see the primary storage device and how much space is being used by various categories of software. Clicking on any of the categories will give you details as to what is in them and give you an opportunity to clean those things up. As you can see in Figure 12.6, this drive is very full and could stand to be cleaned up!

2. In the Storage Management section, click the Cleanup Recommendations link (Figure 12.7).

FIGURE 12.7 Cleanup Recommendations

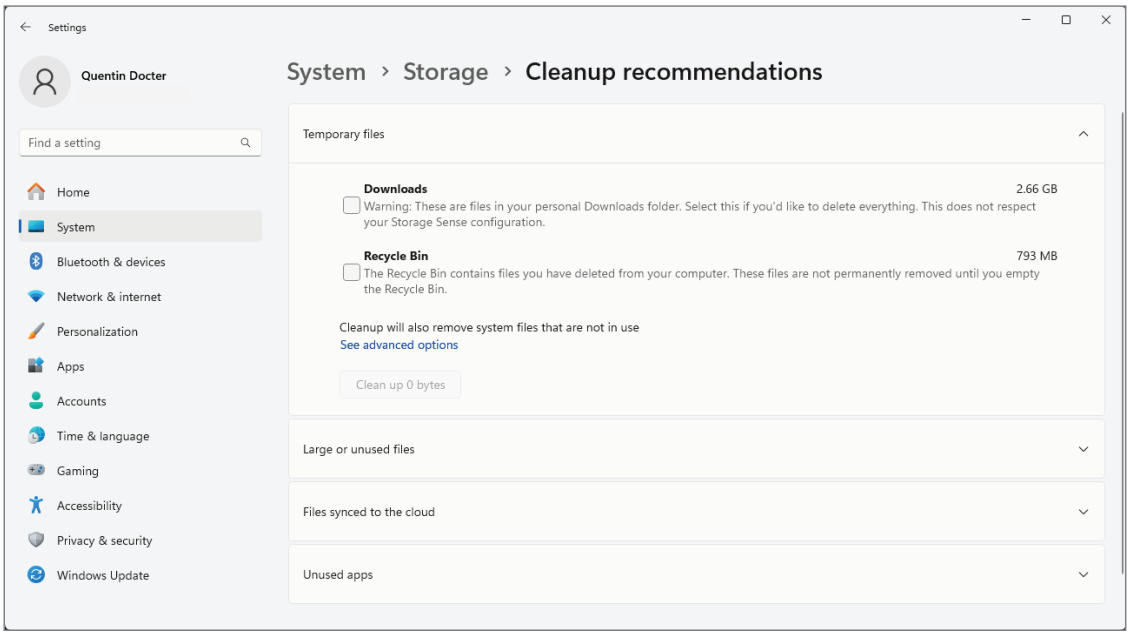
Settings
Quentin Docter
Find a setting
Home
System
Bluetooth & devices
Network & internet
Personalization
Apps
Accounts
Time & language
Gaming
Accessibility
Privacy & security
Windows Update
System › Storage › Cleanup recommendations
Temporary files
Downloads
2.66 GB
Warning: These are files in your personal Downloads folder. Select this if you'd like to delete everything. This does not respect your Storage Sense configuration.
Recycle Bin
793 MB
The Recycle Bin contains files you have deleted from your computer. These files are not permanently removed until you empty the Recycle Bin.
Cleanup will also remove system files that are not in use
See advanced options
Clean up 0 bytes
Large or unused files
Files synced to the cloud
Unused apps

Figure 12.7 shows four categories for potential cleanup, including temporary files, large or unused files, files synced to the cloud, and unused apps. In this specific instance, the cleanup recommendations won't free up much space.

3. In Settings, click the back arrow to get back to the Storage app.
4. In the Storage Management section, click the Advanced Storage Settings link (Figure 12.8).

EXERCISE 12.2 *(continued)*

FIGURE 12.8 Advanced Storage Settings

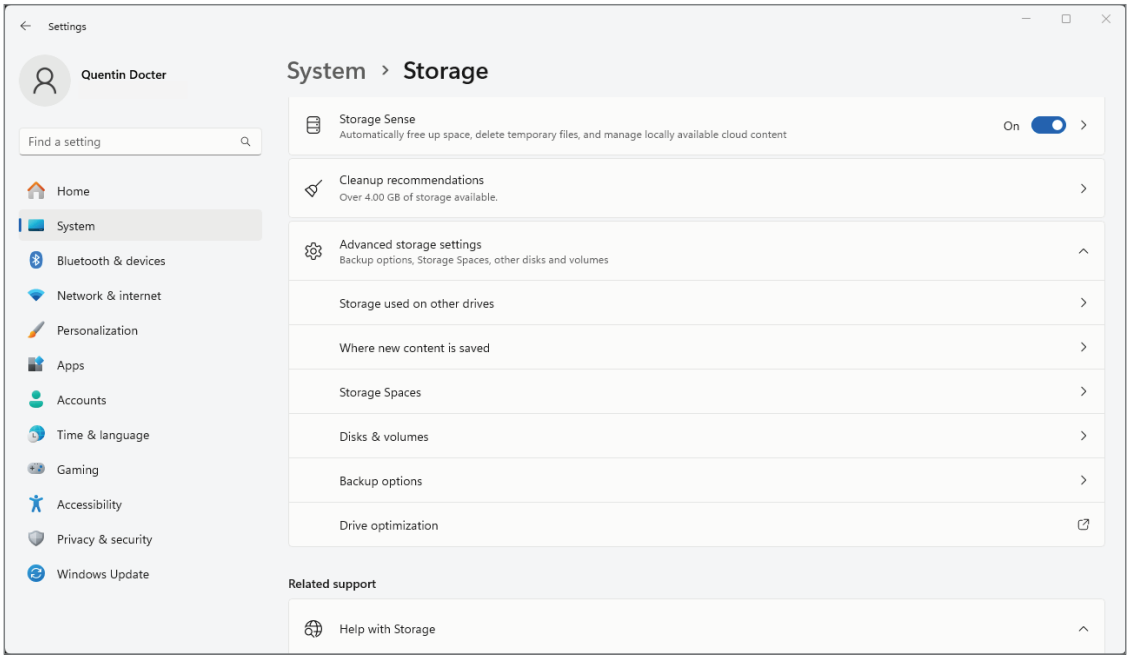

There are several highly functional tools in this section, including the ability to look at all storage drives, determine where new content is saved, set up additional storage spaces, get into Windows backup, and otherwise view and manage storage. We won't go through all of them in this exercise, but we encourage you to click on each option to understand what's there.

5. Click Drive Optimization. It will open the Optimize Drives utility (Figure 12.9). (You can also get to this same utility by typing **Defrag** into the Windows search box—it will pull up the Defragment and Optimize Drives utility for you to open.)

FIGURE 12.9 Optimize Drives

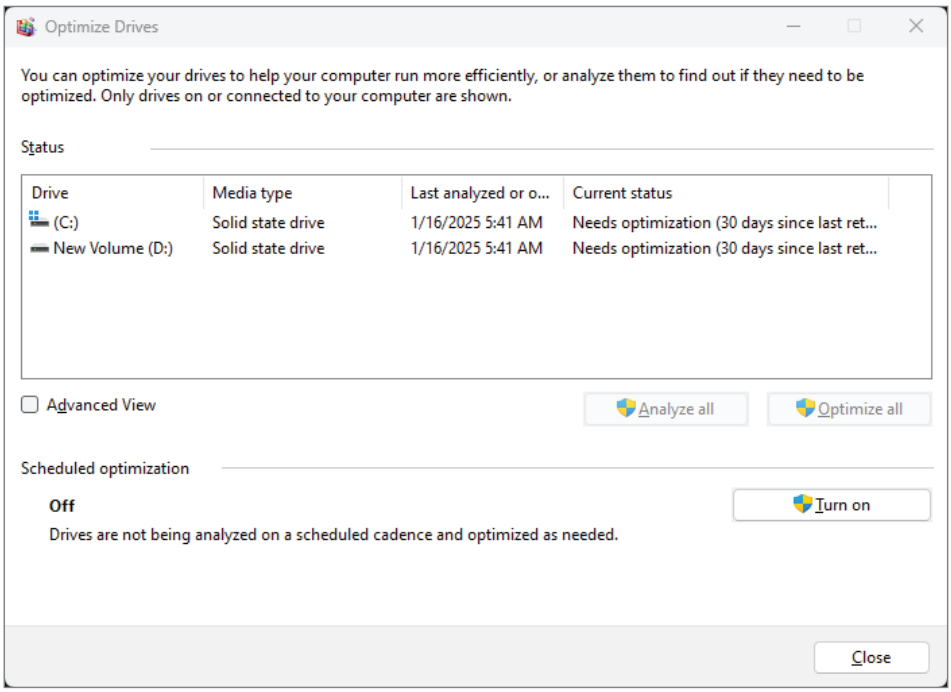
Optimize Drives
You can optimize your drives to help your computer run more efficiently, or analyze them to find out if they need to be optimized. Only drives on or connected to your computer are shown.
Status
Drive
Media type
Last analyzed or o...
Current status
(C:)
Solid state drive
1/16/2025 5:41 AM
Needs optimization (30 days since last ret...
New Volume (D:)
Solid state drive
1/16/2025 5:41 AM
Needs optimization (30 days since last ret...
Advanced View
Analyze all
Optimize all
Scheduled optimization
Off
Turn on
Drives are not being analyzed on a scheduled cadence and optimized as needed.
Close

6. To analyze or optimize a drive, highlight it in the list, and then click Analyze or Optimize. You can also schedule optimization if you'd like.
7. When you are done, close the Settings window.

Input/output operations per second (IOPS), pronounced "eye-ops," is an industry standard for how many reads and writes a storage unit can complete. IOPS is frequently quoted on dedicated storage systems such as NAS and RAID devices, but it's so variable and condition-dependent that its usefulness is debatable. Still, if a device's IOPS steadily declines over time or is no longer fast enough to service the user's (most likely the network's) needs, it could be time to replace the device.

The most popular tool used to measure IOPS is Iometer (`https://sourceforge.net/projects/iometer`); it's open source and it's available for Windows and Linux. Iometer runs simulated disk reads/writes and provides results in a graphical interface (Figure 12.10) and a CSV file. The metrics here might not be the most useful in isolation (e.g., is 3116 I/Os per second good?), but Iometer can be run over time to see if drive performance is declining. For purposes of the A+ exam, don't worry about memorizing any specific metrics or thresholds for IOPS. Instead, understand what it is, know that it can be measured, and know that if performance decreases over time it could indicate an issue with a storage device.

FIGURE 12.10 Iometer test results

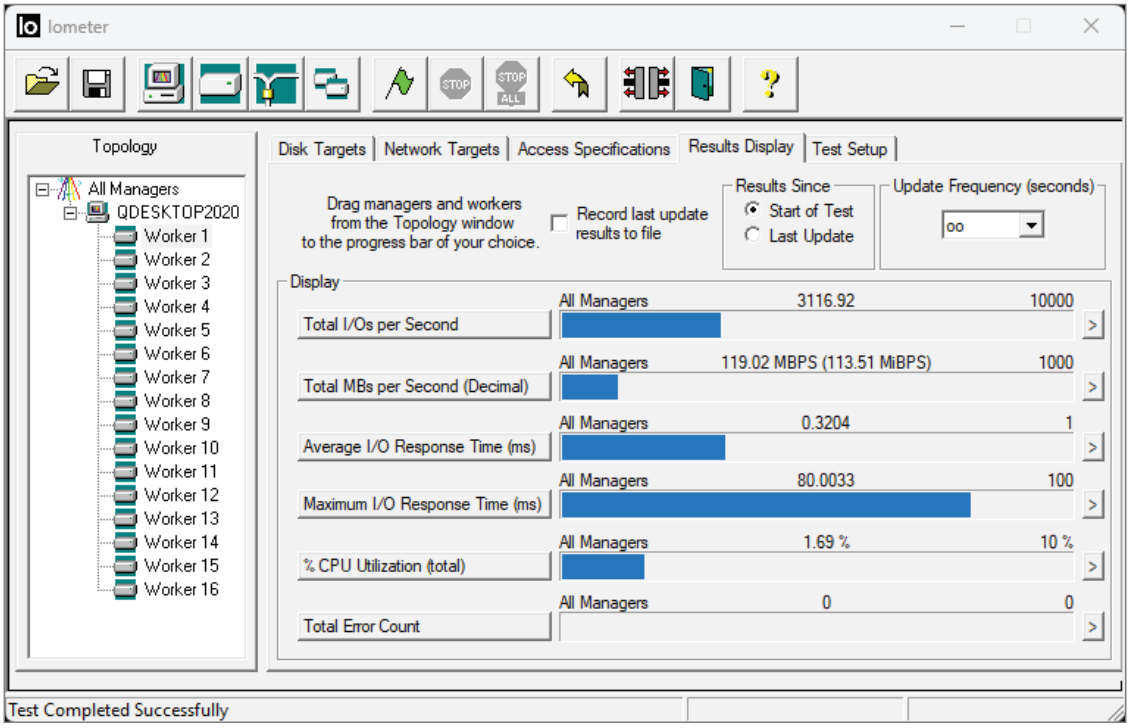

S.M.A.R.T. Diagnostics

As of 2004, nearly every hard drive has been built with *Self-Monitoring, Analysis, and Reporting Technology (S.M.A.R.T.)* software installed. S.M.A.R.T. monitors hard drive reliability and theoretically can warn you in the event of an imminent failure. The idea behind S.M.A.R.T. is great. Who wouldn't want to know when their hard drive was going to fail so they could back up the drive? In practice, though, it seems to help manufacturers locate persistent issues by identifying hard drive design flaws more than it helps end users avoid catastrophic data losses. Helping hard drive manufacturers do a better job isn't a bad thing, but S.M.A.R.T. hasn't enjoyed widespread commercial success with end users. This can largely be attributed to the following three factors:

- Windows OSs don't come with a built-in graphical utility to parse the data.
- The 70 metrics provided by S.M.A.R.T. aren't always easy to understand, and there has been little guidance as to which metric or metrics are most closely associated with impending drive failure.
- Manufacturers have not consistently defined the metrics among themselves; there are no industrywide analysis applications or standards for this technology.

Let's address the three issues in order. First, you can download one of several graphical tools from the Internet if you want to run S.M.A.R.T. diagnostics on a hard drive. Table 12.1 gives you a few options. Each one has a free option, and they all offer a variety of hard drive diagnostic capabilities.

Exam Objective 5.2 specifically mentions S.M.A.R.T. Understand what it is and what it does.

TABLE 12.1 S.M.A.R.T. software utilities

Name	Website
GSmartControl	`https://gsmartcontrol.shaduri.dev`
SpeedFan	`https://almico.com/speedfan.php`
HD Tune Pro	`https://hdtune.com`
CrystalDiskInfo	`https://crystalmark.info/en/software/crystaldiskinfo`

Second, yes, S.M.A.R.T. reports a lot of metrics, not all of which make sense. Figure 12.11 shows the interface of GSmartControl version 2.0.1. On the main page you can see two hard drives, and there's information at the top saying the highlight drive passed the basic health check. Going into more details on the drive gives the output shown in Figure 12.12. There's a lot of info here, and it's hard to know what's really important. For NVMe drives, one metric to look at is Percentage Used (2 percent on this drive). You might look at that and wonder if this is the same almost-full hard drive we showed you in Exercise 12.2, and yes, it is. So why does Percentage Used say only 2 percent? This metric is designed to show the projected endurance for NVMe drives. Remember that SSDs have a finite number of reads and writes before they fail. In this case, about 2 percent of the drive is "used up," meaning it still likely has some longevity left.

If the drive were experiencing failures of some sort, the S.M.A.R.T. diagnostics tool would provide error codes and IDs. In 2014, Google and cloud service provider Backblaze ran large-scale tests to determine which metrics most strongly correlated with drive failure. Their results showed five metrics, which are highlighted in Table 12.2.

FIGURE 12.11 GSmartControl

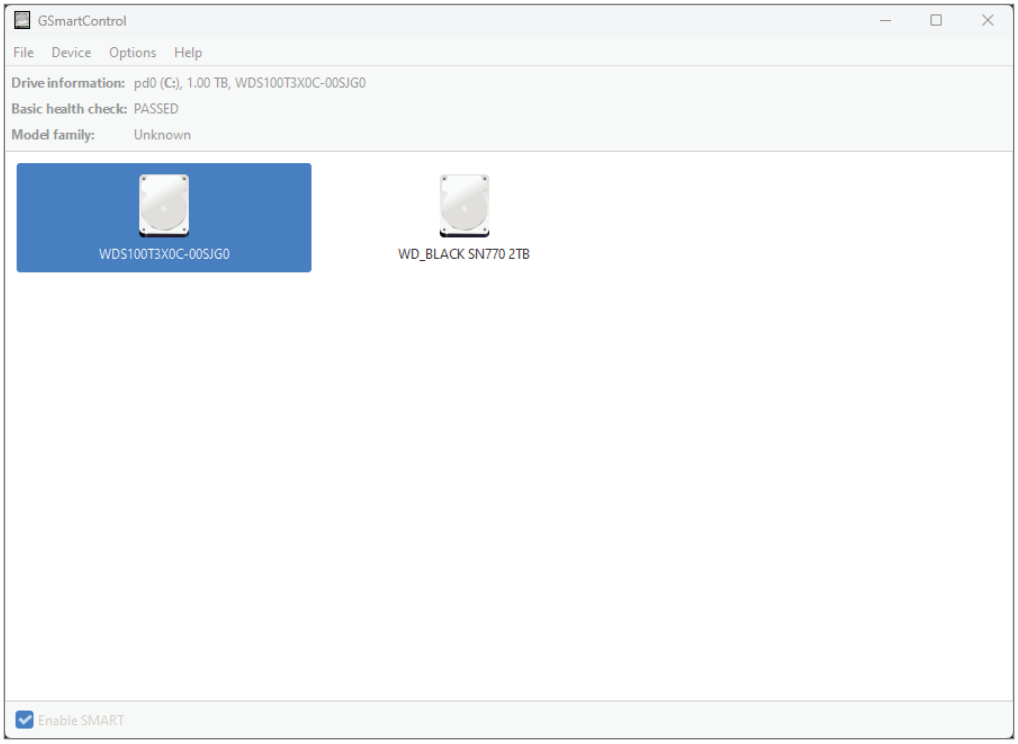

FIGURE 12.12 NVMe information in GSmartControl

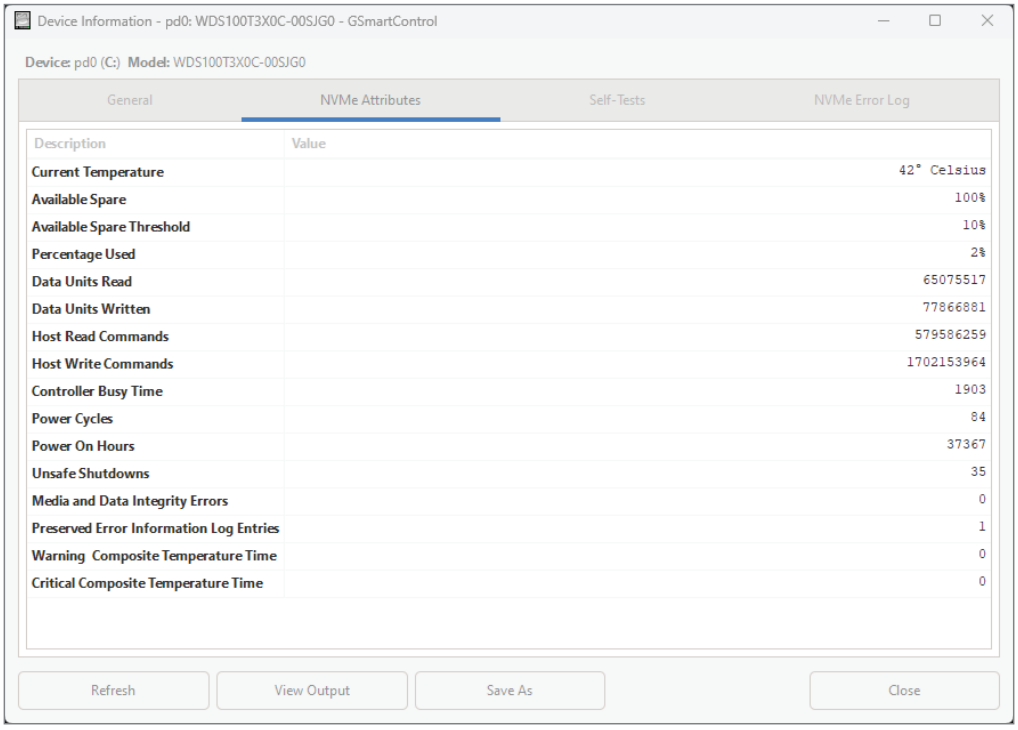

TABLE 12.2 S.M.A.R.T. metrics most correlated with hard drive failure

ID	Attribute name	Description
05	Reallocated sector count	The number of bad sectors that have been found and remapped during read/write processes. Any nonzero number could indicate a problem.
187	Reported uncorrectable errors	The number of errors that could not be recovered using hardware error correction.
188	Command timeout	The number of failed hard drive read/write operations due to disk timeout.
197	Current pending sector count	The number of unstable sectors waiting to be remapped.
198	Uncorrectable sector count	The total number of bad sectors when reading from or writing to a sector.

Interestingly enough, metrics related to higher temperatures or the number of reboots did not correlate to drive failure. The old adage that you should leave your computer running to make the hard drive last longer wasn't verified by the research. In addition, over half of the drives in the study failed without recording a sector error, and over 30 percent of the drives failed with no S.M.A.R.T. error whatsoever. In other words, S.M.A.R.T. might be useful for troubleshooting hard drives, or it might not.

As for the last issue (little consistency between hard drive manufacturers), that's an annoyance but not a critical issue. All it really means is that you can't compare data from one drive manufacturer with that of another. It's likely that if you're running S.M.A.R.T. data on a hard drive, you're primarily concerned with that drive's performance, not how it compares to other hard drives. If you have a situation where you're worried about a drive, you can benchmark its performance and track it over time, or you can just replace it.

RAID Issues

RAID arrays are series of multiple drives working together. You initially learned about them in Volume 1 Chapter 2. If you are using a redundant array of independent (or inexpensive) disks (RAID) system, you have additional challenges to deal with. First, you have more disks, so the chance of having a single failure increases. Second, you more than likely have one or more additional hard disk controllers, so again you introduce more parts that can fail. Third, there will likely be a software component that manages the RAID array.

Exam Objective 5.2 lists two RAID related symptoms: RAID failure and array missing.

Boiling it down, though, dealing with RAID issues is just like dealing with a single hard drive issue, except that you have more parts that make up the single storage unit. If your RAID array is missing or fails, try to narrow down the issue. Is it one disk that's failed, or is the whole system down, indicating a problem with a controller or the software? A single disk failure may cause the RAID array to appear to be missing, and replacing the disk and rebuilding the array should fix the issue. Along with external enclosures, which require a separate connection to the computer, most external RAID systems have status indicators and troubleshooting utilities to help you identify problems. Definitely use those to your advantage.

Finally, the problem could be dependent on the type of RAID you're using. If you are using RAID 0 (disk striping), you actually have more points of failure than a single device, meaning that you're at a greater risk of failure versus using just one hard drive. One drive failure will cause the entire set to fail. RAID 1 (disk mirroring) increases your fault tolerance; if one drive fails, the other has an exact replica of the data. You'll need to replace the failed drive, but unless both drives unexpectedly fail, you shouldn't lose any data. If you're using RAID 5 (disk striping with parity), a minimum of three drives are needed and a single drive failure usually means that your data will be fine, provided that you replace the failed drive. If two or more drives fail, the RAID 5 array will be lost and you will need to fix the array and then restore the data from backup. RAID 6 is like RAID 5 except that it has two parity stripes and requires a minimum of four drives. RAID 10 is a mirrored striped set that requires at least four drives. As long as one drive in each mirrored pair is functional (just like in RAID 1), you shouldn't lose any data.

If your hard drive fails completely and you need to get critical data from it, some third-party companies provide file recovery software and services. These services are generally very expensive. (And you should have been backing up the drive in the first place!)

Optical Drive Issues

Optical drive (Blu-ray, DVD, and CD) problems are normally media-related. Although optical disc technology is pretty reliable, it's not perfect. One factor to consider is the cleanliness of the disc. On many occasions, if a disc is unreadable, cleaning it with an approved cleaner and a lint-free cleaning towel will fix the problem. The next step might be to use a commercially available scratch-removal kit. If that fails, you always have the option to send the disc to a company that specializes in data recovery.

If the operating system doesn't see the drive, start troubleshooting by determining whether the drive is receiving power. If the tray will eject, you can assume the drive has power. Next, check BIOS/UEFI Setup (SATA or PATA drives) to make sure that the drive has been detected. If not, check the primary/secondary jumper on the drive, and make sure that the PATA adapter is set to Auto, CD-ROM, or ATAPI in BIOS/UEFI Setup. Once inside the case, ensure that both the drive and motherboard ends are securely connected and, on a PATA drive, that the ribbon cable is properly aligned with pin 1, the edge that has the red or pink stripe, closest to the power connector.

To play movies, a DVD or Blu-ray drive must have Moving Picture Experts Group (MPEG) decoding capability. This is usually built into the drive, video card, or sound card, but it could require a software decoder. If DVD or Blu-ray data discs will read but not play movies, suspect a problem with the MPEG decoding.

If an optical drive works normally but doesn't perform its special capability (e.g., it won't burn discs), perhaps you need to install software to work with it.

Troubleshooting Networking Problems

As a technician, you are going to be called on to solve a variety of issues, including hardware, software, and networking problems. Networking problems can sometimes be the most tricky to solve, considering that it could be either a software or a hardware problem or a combination of the two causing your connectivity issue.

The first adage for troubleshooting any hardware problem is to check your connections. That holds true for networking as well, but then your troubleshooting will need to go far deeper than that in a hurry. As with troubleshooting anything else, follow a logical procedure and be sure to document your work.

Nearly all the issues tested by CompTIA have something to do with connectivity, which makes sense because that's what networking is all about. Connectivity issues, when not caused by hardware, are generally the result of a messed-up configuration. And because the most common protocol in use today, TCP/IP, has a lot of configuration options, you can imagine how easy it is to configure something incorrectly.

In the following sections, we'll look at connectivity issues and how to resolve them. But first, we will introduce you to several hardware tools and software commands you should know about, as they are essential to network troubleshooting.

Exam Objective 5.5, "Given a scenario, troubleshoot network issues," combines wired and wireless networking. We'll cover them together and note when specific differences exist.

Using Network Troubleshooting Tools

Networks are so specialized that they have their own set of troubleshooting tools. This includes hardware and software tools. Knowing how to use them can make the difference between a quick fix and a lengthy and frustrating troubleshooting session. The CompTIA A+ exams will test you on both hardware and software tools, as well as give you troubleshooting scenarios that require your knowledge of how they work. Hardware tools are covered in Exam 220-1201, whereas the software commands are part of Exam 220-1202, Objective 1.5. Because both are critical to resolving network issues, we will cover them all here. We'll then review the software tools again in Volume 2 Chapter 3.

Hardware and Cabling Tools

We covered several different types of cables and their properties in Volume 1 Chapter 6, "Networking Fundamentals." Here, we will look at some tools that can be used to make or test network cables, as well as a few tools to troubleshoot network connectivity issues.

Exam Objective 2.8, "Explain networking tools and their purposes," lists eight tools you should know: crimper, cable stripper, Wi-Fi analyzer, toner probe, punchdown tool, cable tester, loopback plug, and network tap.

Multimeter Multimeters are versatile electronic measuring tools. A *multimeter* can measure voltage, current, and resistance on a wire. There are many different types and qualities on the market, everywhere from economical $10 versions to ones that cost several thousand dollars. Figure 12.13 shows a basic multimeter. Note that of the tools we cover, the multimeter is the only one that's not on the exam objectives. In real life, they are very useful.

FIGURE 12.13 A multimeter

Crimpers and Cable Strippers A *crimper* is a very handy tool for helping you put connectors on the end of a cable. Most crimpers are combination tools that strip and snip

wires and crimp the connector onto the end (see Figure 12.14). You can also buy tools that are just *cable strippers* and cable snippers (Figure 12.15), but usually the point of cutting and stripping a cable is to put a connector on the end of it, so why not just get a crimper that does it all?

FIGURE 12.14 A UTP crimper and connectors

FIGURE 12.15 Cable stripper

Wi-Fi Analyzer One hardware tool that we didn't talk much about earlier is the *Wi-Fi analyzer*. Specific tools for locating Wi-Fi networks and analyzing their traffic are indispensable today. A wireless locator or a Wi-Fi analyzer can be either a handheld hardware device, such as the one shown in Figure 12.16, or specialized software that is installed on a laptop and the purpose of which is to detect and analyze Wi-Fi signals. Anyone interested in wardriving (driving around in a car looking for unsecured wireless networks to hack into) will definitely have one of these, but they're also handy for locating wireless hot spots.

FIGURE 12.16 Handheld Wi-Fi analyzer

Toner Probe If you need to trace a wire in a wall from one location to another, a *toner probe* (also called a tone generator and probe) is the tool for you. Shown in Figure 12.17, it consists of two pieces: a tone generator and a probe. Because it's so good at tracking, you will sometimes hear this referred to as a "fox and hound."

FIGURE 12.17 Toner probe

To use a toner probe, attach one end to one end of the cable, such as the end at the computer. Then go to the patch panel with the other end of the probe to locate the cable. These are lifesavers when the cables are not properly labeled.

Punchdown Tool If you're working on a larger network installation, you might use a *punchdown tool*. It's not a testing tool but one that allows you to connect (i.e., punch down) the exposed ends of a wire into wiring harnesses, such as a 110 block (often used in connectivity closets to help simplify the tangled mess of cables). Figure 12.18 shows the tool and its changeable bit.

FIGURE 12.18 Punchdown tool

Cable Tester A *cable tester* is an indispensable tool for any network technician. Usually, you would use a cable tester before you install a cable to make sure it works. Of course, you can test it after it's been run as well. A decent cable tester will tell you the type of cable, and more elaborate models have connectors for multiple types of cables. Figure 12.19 shows a TRENDnet cable tester.

Loopback Plug A *loopback plug* is for testing the ability of a network adapter to send and receive. The plug gets plugged into the NIC, and then a loopback test is performed using troubleshooting software. You can then tell whether or not the card is working properly. Figure 12.20 shows an Ethernet loopback plug, and they are made for fiber-optic NICs as well.

Network Tap The final tool we'll cover is indispensable for network troubleshooting. A *network tap* (*tap* is an acronym for test access port) is a hardware device that creates a copy of network traffic for use by monitoring devices without disrupting normal traffic. Taps have great versatility because they can be moved from one location to another on a network, allowing an administrator to focus on an area that's believed to have problems. They can also detect potential network intrusions, making them helpers for network security as well. In Figure 12.21, you can see where the network cables plug into; the monitor plugs into the back side.

FIGURE 12.19 Cable tester

FIGURE 12.20 Ethernet loopback plug

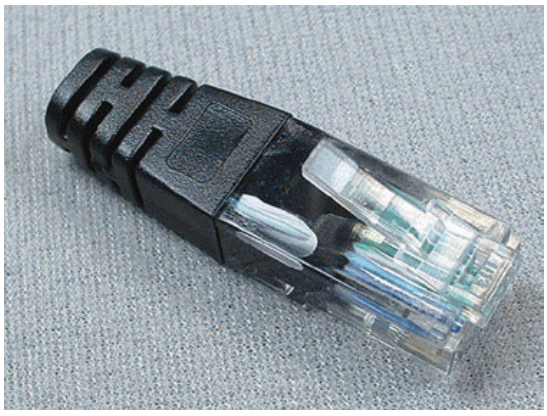

FIGURE 12.21 Network tap

Software Commands

Troubleshooting networks often involves using a combination of hardware tools and software commands. Usually, the software commands are easier to deal with because you don't need to dig around physically in a mess of wires to figure out what's going on. The downside to the software commands is that there can be a number of options that you need to memorize. In the following sections, we'll cover the following networking command-line tools, which are all helpful utilities: `ipconfig/ip`, `ping`, `hostname`, `tracert`, `netstat`, `net`, and `nslookup`.

Again, these software commands are not specifically listed in the A+ 220-1201 exam objectives, but they are in 220-1202 Exam Objective 1.5. We include them here because we don't want you to get caught off guard if you get asked troubleshooting scenarios on the 220-1201 exam that require knowledge of these tools.

ipconfig Command

With Windows-based operating systems, you can determine the network settings on the client's network interface cards, as well as any that a DHCP server has leased to your computer, by typing the following at a command prompt: **`ipconfig /all`**.

`ipconfig /all` also gives you full details on the duration of your current lease. You can verify whether a DHCP client has connectivity to a DHCP server by releasing the client's IP address and then attempting to lease an IP address. You can conduct this test by typing the following sequence of commands from the DHCP client at a command prompt:

```
ipconfig /release
ipconfig /renew
```

`ipconfig` is one of the first tools to use when experiencing problems accessing resources because it will show you whether an address has been issued to the machine. If the address displayed falls within the `169.254.x.x` category, this means that the client was unable to reach the DHCP server and has defaulted to Automatic Private IP Addressing (APIPA), which will prevent the network card from communicating outside its subnet, if not altogether. Table 12.3 lists useful switches for `ipconfig`.

TABLE 12.3 `ipconfig` switches

Switch	Purpose
`/all`	Shows full configuration information
`/release`	Releases the IP address if you are getting addresses from a Dynamic Host Configuration Protocol (DHCP) server
`/release6`	Releases the IPv6 addresses
`/renew`	Obtains a new IP address from a DHCP server
`/renew6`	Obtains a new IPv6 address from a DHCP server
`/flushdns`	Flushes the Domain Name System (DNS) server's name resolver cache

Whenever you are dealing with commands at the command prompt, you can always use the `/?` switch to see all available options for that command.

In the Linux, UNIX, and macOS worlds, a utility similar to `ipconfig` is `ip`.

Figure 12.22 shows the output from `ipconfig`, and Figure 12.23 shows the output from `ipconfig /all` for the same network adapter.

FIGURE 12.22 `ipconfig` output

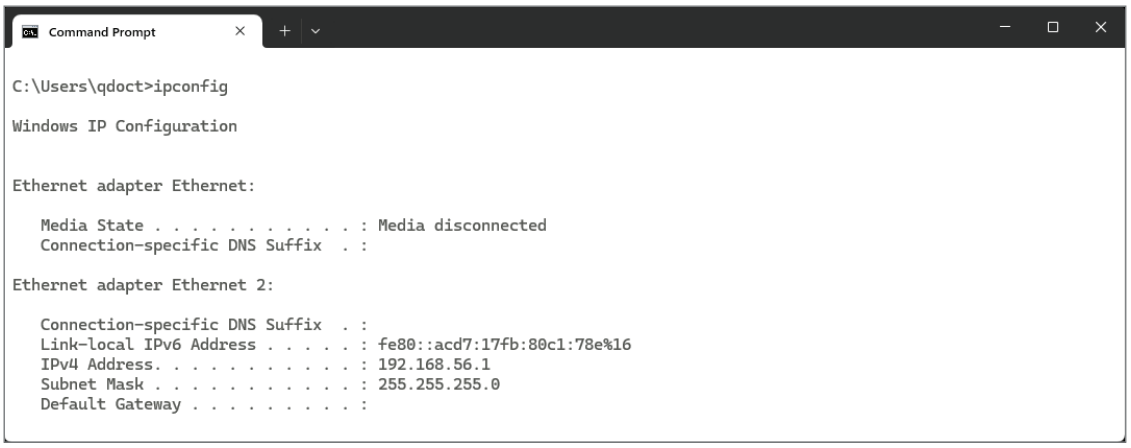
Command Prompt
C:\Users\qdoct>ipconfig
Windows IP Configuration
Ethernet adapter Ethernet:
Media State . . . . . : Media disconnected
Connection-specific DNS Suffix . :
Ethernet adapter Ethernet 2:
Connection-specific DNS Suffix . :
Link-local IPv6 Address . . . . . : fe80::acd7:17fb:80c1:78e%16
IPv4 Address. . . . . : 192.168.56.1
Subnet Mask . . . . . : 255.255.255.0
Default Gateway . . . . . :

FIGURE 12.23 `ipconfig /all` output

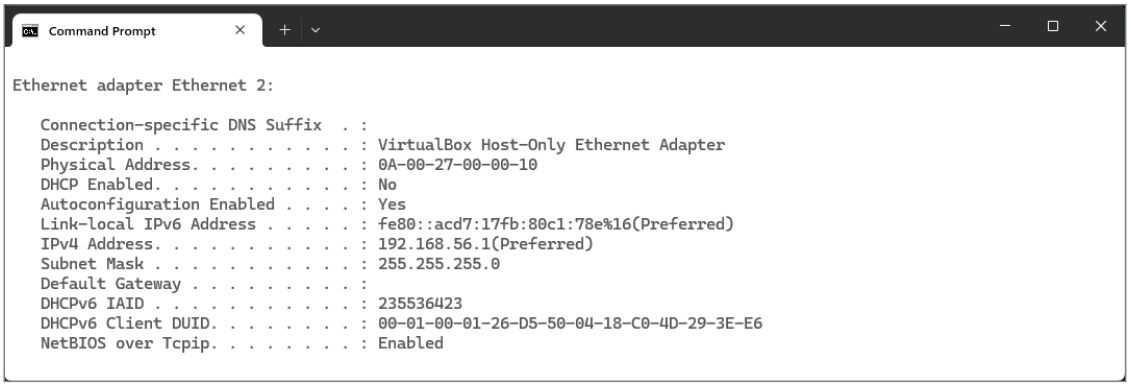
Command Prompt
Ethernet adapter Ethernet 2:
Connection-specific DNS Suffix . :
Description . . . . . : VirtualBox Host-Only Ethernet Adapter
Physical Address. . . . . : 0A-00-27-00-00-10
DHCP Enabled. . . . . : No
Autoconfiguration Enabled . . . . : Yes
Link-local IPv6 Address . . . . . : fe80::acd7:17fb:80c1:78e%16(Preferred)
IPv4 Address. . . . . : 192.168.56.1(Preferred)
Subnet Mask . . . . . : 255.255.255.0
Default Gateway . . . . . :
DHCPv6 IAID . . . . . : 235536423
DHCPv6 Client DUID. . . . . : 00-01-00-01-26-D5-50-04-18-C0-4D-29-3E-E6
NetBIOS over Tcpip. . . . . : Enabled

In Exercise 12.3 , you will renew an IP address on a Windows 11 system within the graphical interface.

EXERCISE 12.3

Renewing an IP Address in Windows 11

In the Windows interface, getting a new IP address means disabling the current connection and then re-enabling it. This is done through the connection's properties. You can get there in Settings in a few different ways, and we will show you one way here.

1. Click Start ➢ Settings ➢ Network & Internet to open the Network & Internet app.
2. At the bottom of the screen, click Advanced Network Settings.

EXERCISE 12.3 *(continued)*

3. On the Advanced Network Settings screen, the Network Adapters section shows the configured network adapters. We will use this area in a few steps. For now, in the More Settings section, click Hardware And Connection Properties (Figure 12.24).

FIGURE 12.24 Hardware And Connection Properties

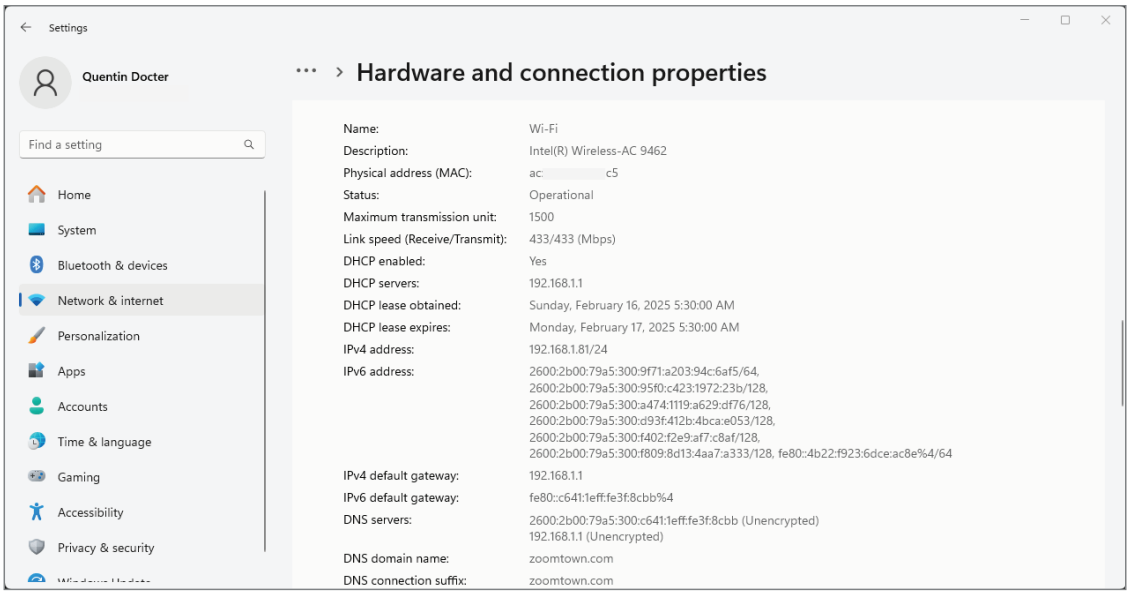

4. Note the time the lease was obtained.
5. Click the Back arrow in the top-left corner of the screen.
6. On the Advanced Network Settings page, click the Disable button next to the network connection.

   The list of network adapters might completely disappear. If so, click the Back arrow in the upper-left corner, and then click Advanced Network Settings. The list of network adapters should reappear.
7. Next to the network connection name, click Enable. It should take a second or two to re-enable the connection.
8. Click Hardware And Connection Properties.
9. Note the time the lease was obtained. It should be very recent.

While Windows allows you to renew IP addresses and troubleshoot connection problems from the Network & Internet app, some administrators still prefer the reliability of a command-line interface. Exercise 12.4 shows you how to perform a similar action using the command line.

EXERCISE 12.4

Renewing an IP Address from the Command Line

This exercise assumes that you are using Windows and dynamic IP assignments from a DHCP server.

1. Open a command prompt (type **cmd** into the Windows search box and press Enter).
2. Type **ipconfig** and view the abbreviated list of information.
3. Type **ipconfig /all** to see the full list. Notice the date and time on the lease for the IP address.
4. Type **ipconfig /renew** and press Enter.
5. Type **ipconfig /all** again. The date and time on the lease for the IP address should be the current date and time.
6. Close the command-prompt window by typing **exit** and pressing Enter.

ping Command

The `ping` command is one of the most useful commands in the TCP/IP protocol. It sends a series of packets to another system, which in turn sends back a response. This utility can be extremely useful for troubleshooting problems with remote hosts. Pings are also called ICMP echo requests/replies because they use Internet Control Message Protocol (ICMP).

The `ping` command indicates whether the host can be reached and how long it took for the host to send a return packet. Across WAN links, the time value will be much larger than across healthy LAN links.

The syntax for `ping` is `ping` *`hostname`* or `ping` *`IP address`*. Figure 12.25 shows what a ping should look like.

As you can see, by pinging with the hostname, we found the host's IP address thanks to DNS. The time is how long in milliseconds it took to receive the response. On a LAN, you want this to be 10 milliseconds (ms) or less, but 25ms for an Internet ping is pretty good.

The `ping` command has several options, which you can see by typing **`ping /?`** at the command prompt. Table 12.4 lists some useful options.

FIGURE 12.25 A successful ping

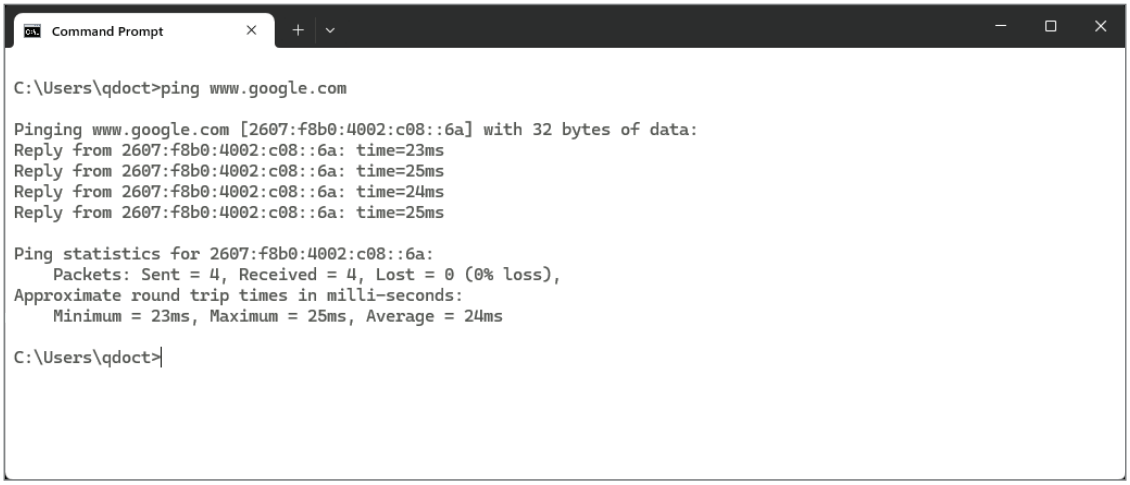
Command Prompt
C:\Users\qdoct>ping www.google.com
Pinging www.google.com [2607:f8b0:4002:c08::6a] with 32 bytes of data:
Reply from 2607:f8b0:4002:c08::6a: time=23ms
Reply from 2607:f8b0:4002:c08::6a: time=25ms
Reply from 2607:f8b0:4002:c08::6a: time=24ms
Reply from 2607:f8b0:4002:c08::6a: time=25ms
Ping statistics for 2607:f8b0:4002:c08::6a:
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 23ms, Maximum = 25ms, Average = 24ms
C:\Users\qdoct>

TABLE 12.4 `ping` options

Option	Function
`-t`	Persistent ping. Will ping the remote host until stopped by the client (by using Ctrl+C).
`-n` *count*	Specifies the number of echo requests to send.
`-l` *size*	Specifies the packet size to send.
`ping -4` / `ping -6`	Uses either the IPv4 or IPv6 network, respectively.

Some webmasters have configured their routers to block pings in order to avoid problems such as someone trying to eat up bandwidth with a *ping of death* (sending a persistent ping with a huge buffer to overwhelm the recipient). Therefore, if you ping a website, it's possible that you won't get a response even though the site is functional.

hostname Command

The `hostname` command is a very simple one. It returns the name of the host computer on which its executed. Figure 12.26 shows you the output.

FIGURE 12.26 `hostname` output

netstat Command

The `netstat` command is used to check out the inbound and outbound TCP/IP connections on your machine. It can also be used to view packet statistics, such as how many packets have been sent and received and the number of errors.

When used without any options, the `netstat` command produces output similar to that shown in Figure 12.27, which shows all the outbound TCP/IP connections.

FIGURE 12.27 `netstat` output

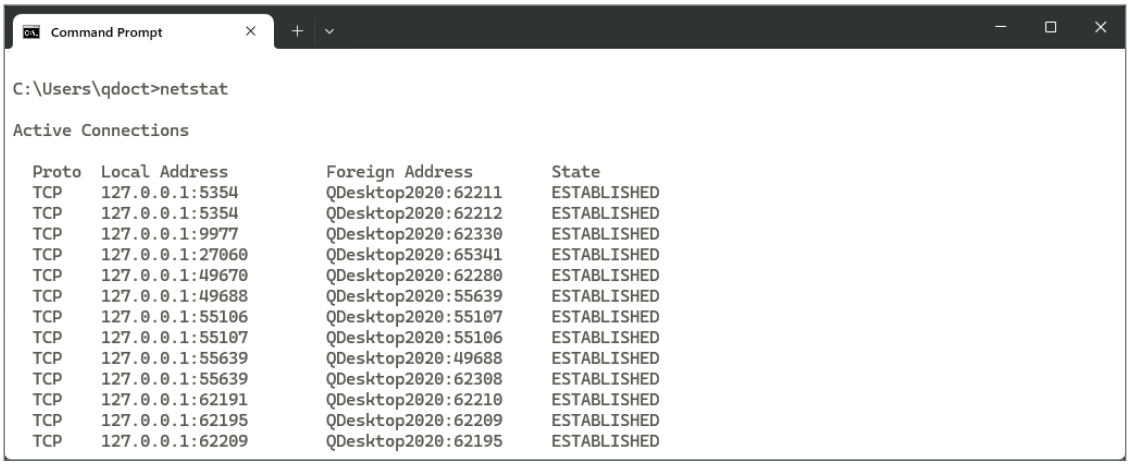

There are several useful command-line options for `netstat`, as shown in Table 12.5.

TABLE 12.5 `netstat` options

Option	Function
`-a`	Displays all connections and listening ports.
`-b`	Displays the executable involved in creating each connection or listening port. In some cases, well-known executables host multiple independent components, and in these cases the sequence of components involved in creating the connection or listening port is displayed. In this case, the executable name is in brackets at the bottom. At the top are the components it called, in sequence, until TCP/IP was reached. Note that this option can be time consuming, and it will fail unless you have sufficient permissions.
`-e`	Displays Ethernet statistics. This may be combined with the `-s` option.
`-f`	Displays fully qualified domain names (FQDNs) for foreign addresses.
`-n`	Displays addresses and port numbers in numerical form.
`-o`	Displays the owning process ID associated with each connection.
`-p` *`proto`*	Shows connections for the protocol specified by *`proto`*; *`proto`* may be any of the following: TCP, UDP, TCPv6, or UDPv6. If `netstat` is used with the `-s` option to display per-protocol statistics, *`proto`* may be IP, IPv6, ICMP, ICMPv6, TCP, TCPv6, UDP, or UDPv6.
`-r`	Displays the routing table.
`-s`	Displays per-protocol statistics. By default, statistics are shown for IP, IPv6, ICMP, ICMPv6, TCP, TCPv6, UDP, and UDPv6; the `-p` option may be used to specify a subset of the default.

nslookup Command

One of the key things that must take place to use TCP/IP effectively is that a hostname must resolve to an IP address—an action usually performed by a DNS server.

`nslookup` is a command that enables you to verify entries on a DNS server. You can use the `nslookup` command in two modes: interactive and noninteractive. In interactive mode, you start a session with the DNS server in which you can make several requests. In noninteractive mode, you specify a command that makes a single query of the DNS server. If you want to make another query, you must type another noninteractive command.

To start `nslookup` in interactive mode (which is what most admins use because it allows them to make multiple requests without typing `nslookup` several times), type **`nslookup`** at the command prompt and press Enter. You will receive a greater than prompt (>) and you can then type the command that you want to run. You can also type **`help`** or **`?`** to bring up

the list of possible commands, as shown in Figure 12.28. To exit `nslookup` and return to a command prompt, type **`exit`** and press Enter.

FIGURE 12.28 Starting `nslookup` and using `help`

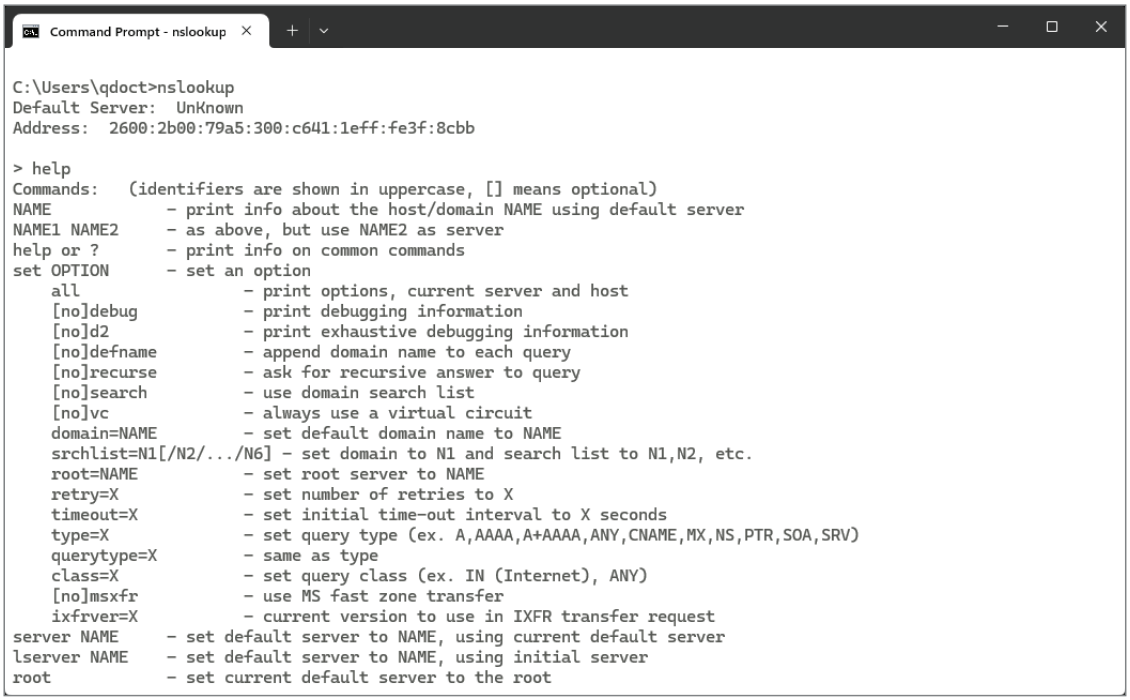
Command Prompt - nslookup
C:\Users\qdoct>nslookup
Default Server: UnKnown
Address: 2600:2b00:79a5:300:c641:1eff:fe3f:8cbb
> help
Commands: (identifiers are shown in uppercase, [] means optional)
NAME   - print info about the host/domain NAME using default server
NAME1 NAME2   - as above, but use NAME2 as server
help or ?   - print info on common commands
set OPTION   - set an option
    all   - print options, current server and host
    [no]debug   - print debugging information
    [no]d2   - print exhaustive debugging information
    [no]defname   - append domain name to each query
    [no]recurse   - ask for recursive answer to query
    [no]search   - use domain search list
    [no]vc   - always use a virtual circuit
    domain=NAME   - set default domain name to NAME
    srchlist=N1[/N2/.../N6] - set domain to N1 and search list to N1,N2, etc.
    root=NAME   - set root server to NAME
    retry=X   - set number of retries to X
    timeout=X   - set initial time-out interval to X seconds
    type=X   - set query type (ex. A,AAAA,A+AAAA,ANY,CNAME,MX,NS,PTR,SOA,SRV)
    querytype=X   - same as type
    class=X   - set query class (ex. IN (Internet), ANY)
    [no]msxfr   - use MS fast zone transfer
    ixfrver=X   - current version to use in IXFR transfer request
server NAME   - set default server to NAME, using current default server
lserver NAME   - set default server to NAME, using initial server
root   - set current default server to the root

To run `nslookup` in noninteractive mode, you would use the `nslookup` command option you want to run at the command prompt—for example, **`nslookup /set timeout=<3>`** or **`nslookup /view:domain`**.

net Command

Depending on the version of Windows you are using, `net` can be one of the most powerful commands at your disposal. All Windows versions include a `net` command, but its capabilities differ based on whether it is used on a server or workstation and the version of the operating system.

While always command line–based, `net` allows you to do almost anything that you want with the operating system. Table 12.6 shows common `net` switches.

TABLE 12.6 `net` switches

Switch	Purpose
`net accounts`	To set account options (password age, length, and so on)
`net computer`	To add and delete computer accounts

TABLE 12.6 net switches *(continued)*

Switch	Purpose
`net config`	To see network-related configuration
`net continue`, `net pause`, `net start`, `net statistics`, and `net stop`	To control services
`net file`	To close open files
`net group` and `net localgroup`	To create, delete, and change groups
`net help`	To see general help
`net helpmsg`	To see specific message help
`net name`	To see the name of the current machine and user
`net print`	To interact with print queues and print jobs
`net send`	To send a message to user(s)
`net session`	To see session statistics
`net share`	To create a share
`net time`	To set the time to that of another computer
`net use`	To connect to a share
`net user`	To add, delete, and see information about a user
`net view`	To see available resources

These commands are invaluable troubleshooting aids when you cannot get the graphical interface to display properly. You can also use them when interacting with hidden (`$`) and administrative shares that do not appear within the graphical interface.

The `net` command used with the `share` parameter enables you to create shares from the command prompt, using this syntax:

```
net share <share_name>=<drive_letter>:<path>
```

To share the C:\EVAN directory as SALES, you would use the following command:

```
net share sales=c:\evan
```

You can use other parameters with `net share` to set other options. Table 12.7 summarizes the most commonly used parameters, and Exercise 12.5 will give you some experience with the `net share` command.

TABLE 12.7 `net share` parameters

Parameter	Purpose
/delete	To stop sharing a folder
/remark	To add a comment for browsers
/unlimited	To set the user limit to Maximum Allowed
/users	To set a specific user limit

EXERCISE 12.5

Using the `net share` Command in Windows

In order to create new shared drives with `net share`, you must have administrative privileges at the command prompt. Opening the command prompt normally does not confer administrative privileges. In this exercise, we will open what's called an elevated command prompt, or administrative command prompt.

1. Type **cmd** in the Windows search box and notice that the Command Prompt appears under Best Match.
2. On the right side, under the Command Prompt icon, there is an option to Run As Administrator. Click that. (You can also right-click the app icon and then choose Run As Administrator.)
3. A User Account Control prompt will appear asking if you want to allow the app to make changes to your device. Click Yes. The elevated command prompt will open. Notice that in the title bar of the command prompt, it says Administrator: Command Prompt.
4. To see what is currently shared on your computer, type **net share** and press Enter.

   You will see output similar to what's shown in Figure 12.29. Notice that the C:\ and D:\ drives are shared by default. This is normal, and they are hidden shares (the $ at the end of the share name makes it hidden), meaning that other users on the network won't see them.

EXERCISE 12.5 *(continued)*

FIGURE 12.29 `net share` command

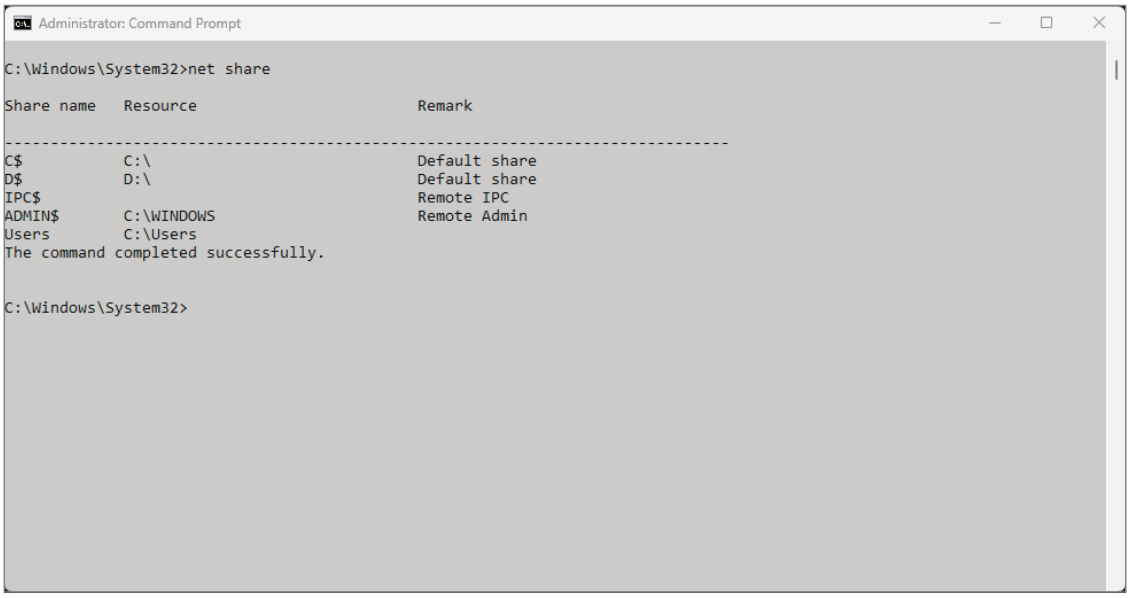

5. You will now create a `C:\TEMP` folder to share.
6. Type **`cd\`** and press Enter to get to the `c:\` prompt.
7. Type **`md temp`** and press Enter.
8. Type **`dir`** and you should see the `temp` folder (see Figure 12.30).

FIGURE 12.30 Administrator Command Prompt

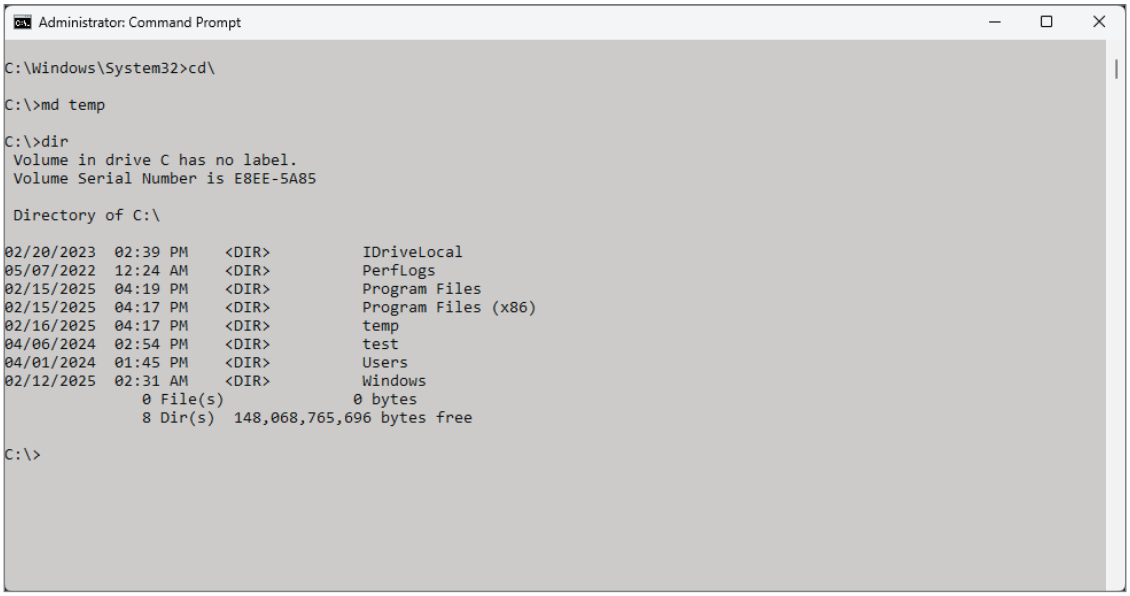

9. Share the `C:\TEMP` folder with the share name `files4u`. To do this, type **`net share files4u=c:\temp`** and press Enter.

   The system should respond with `files4u was shared successfully`.

10. To verify that it was shared, type **`net share`** and press Enter.

    You can see the new share in Figure 12.31. Users on the network can now access your drive as `\\`*`computername`*`\files4u`.

FIGURE 12.31 Sharing a folder

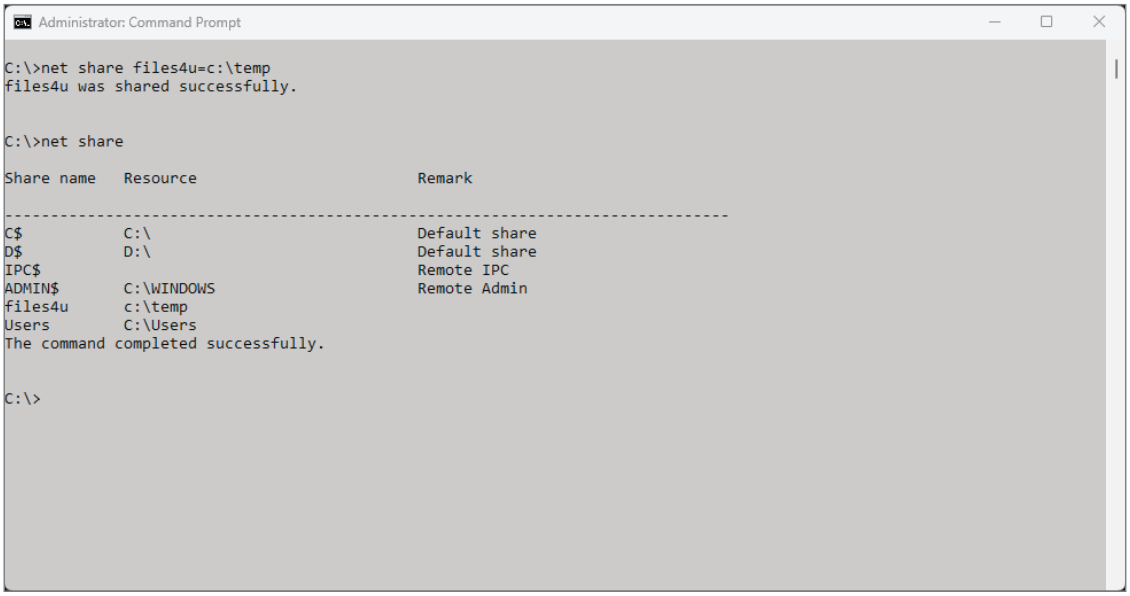

11. (optional) To stop sharing `files4u`, type **`net share files4u /delete`** at the command prompt and press Enter.

12. (optional) To verify that the folder is no longer shared, type **`net share`** and press Enter.

13. Close the command prompt.

The `net /?` command is basically a catch-all help request. It will instruct you to use the `net` command in which you are interested for more information.

tracert Command

`tracert` (trace route) is a Windows-based command-line utility that enables you to verify the route to a remote host. Execute the command `tracert` *`hostname`*, where *`hostname`* is the computer name or IP address of the computer whose route you want to trace. `tracert` returns

the different IP addresses the packet was routed through to reach the final destination. The results also include the number of hops needed to reach the destination. If you execute the `tracert` command without any options, you see a help file that describes all the `tracert` switches.

The `traceroute` command can be used in macOS and Linux/UNIX; it performs the same task as `tracert`.

This utility determines the intermediary steps involved in communicating with another IP host. It provides a road map of all the routing an IP packet takes to get from host A to host B.

Timing information from `tracert` can be useful for detecting a malfunctioning or overloaded router. Figure 12.32 shows the output from `tracert`. In addition to `tracert`, many graphical third-party network-tracing utilities are available on the market.

FIGURE 12.32 `tracert` output

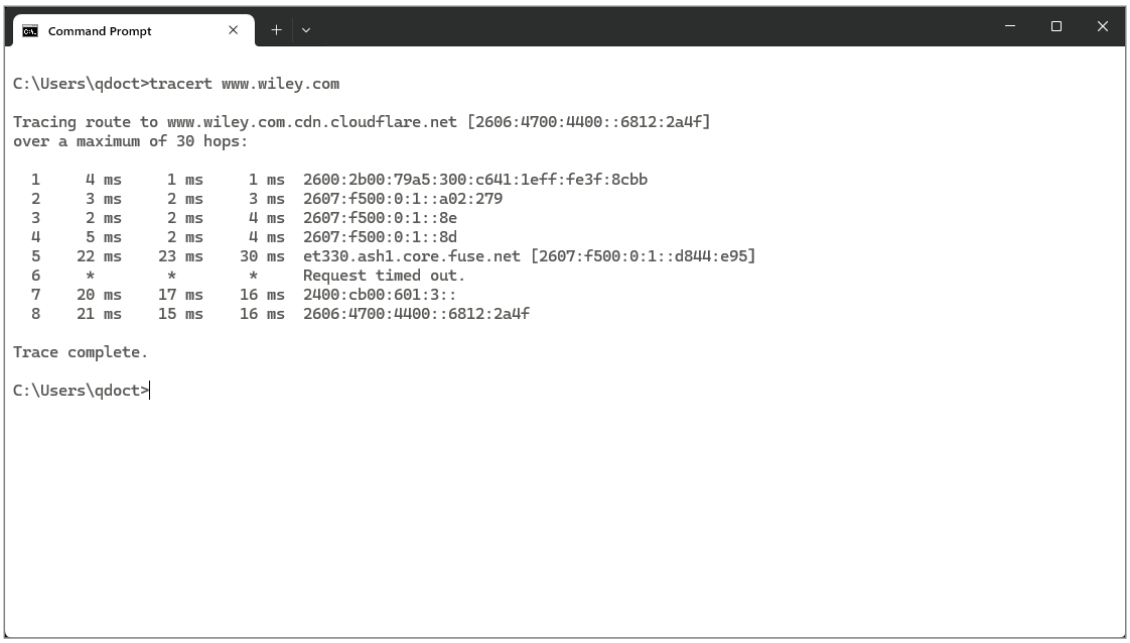
Command Prompt
C:\Users\qdoct>tracert www.wiley.com
Tracing route to www.wiley.com.cdn.cloudflare.net [2606:4700:4400::6812:2a4f]
over a maximum of 30 hops:
1 4 ms 1 ms 1 ms 2600:2b00:79a5:300:c641:1eff:fe3f:8cbb
2 3 ms 2 ms 3 ms 2607:f500:0:1::a02:279
3 2 ms 2 ms 4 ms 2607:f500:0:1::8e
4 5 ms 2 ms 4 ms 2607:f500:0:1::8d
5 22 ms 23 ms 30 ms et330.ash1.core.fuse.net [2607:f500:0:1::d844:e95]
6 * * * Request timed out.
7 20 ms 17 ms 16 ms 2400:cb00:601:3::
8 21 ms 15 ms 16 ms 2606:4700:4400::6812:2a4f
Trace complete.
C:\Users\qdoct>

pathping Command

`pathping` (path ping) combines the best of both worlds from `ping` and `tracert` and is a favorite command for many net admins. It acts much like a ping, but it also traces the route to the destination and shows where, if anywhere, packet loss occurs between the sending computer and the remote host. Said differently, `pathping` first traces the route to the destination host, and then it pings each node between itself and the destination. Figure 12.33 shows the output. A similar command for Linux-based computers is `mtr`.

FIGURE 12.33 `pathping` output

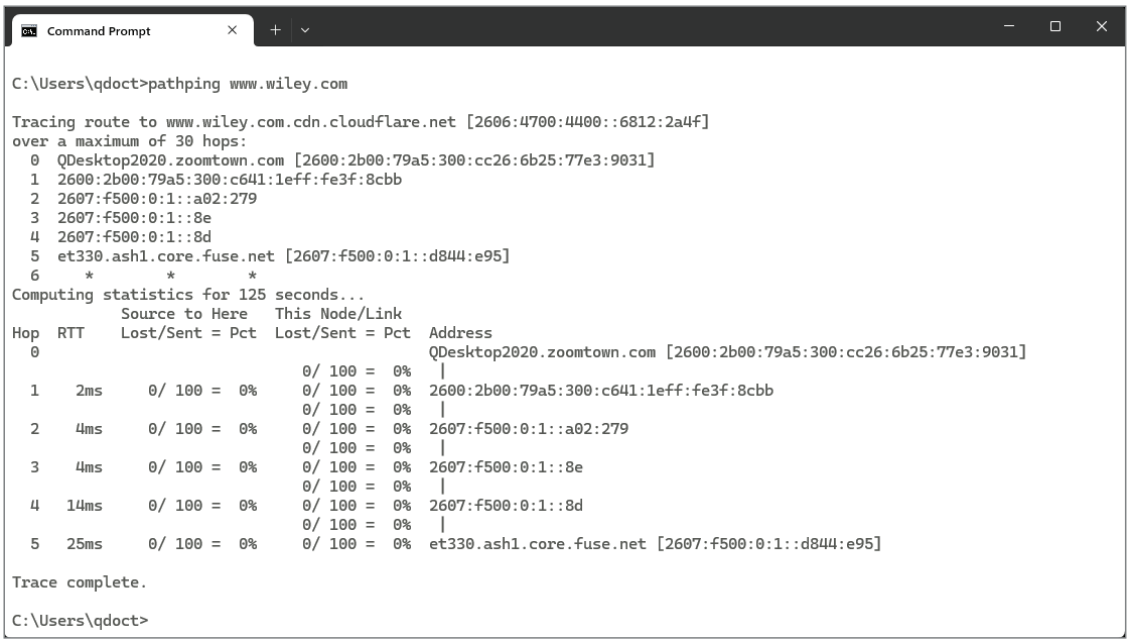
Command Prompt
C:\Users\qdoct>pathping www.wiley.com
Tracing route to www.wiley.com.cdn.cloudflare.net [2606:4700:4400::6812:2a4f]
over a maximum of 30 hops:
0 QDesktop2020.zoomtown.com [2600:2b00:79a5:300:cc26:6b25:77e3:9031]
1 2600:2b00:79a5:300:c641:1eff:fe3f:8cbb
2 2607:f500:0:1::a02:279
3 2607:f500:0:1::8e
4 2607:f500:0:1::8d
5 et330.ash1.core.fuse.net [2607:f500:0:1::d844:e95]
6 * * *
Computing statistics for 125 seconds...
Source to Here This Node/Link
Hop RTT Lost/Sent = Pct Lost/Sent = Pct Address
0 QDesktop2020.zoomtown.com [2600:2b00:79a5:300:cc26:6b25:77e3:9031]
0/ 100 = 0% |
1 2ms 0/ 100 = 0% 0/ 100 = 0% 2600:2b00:79a5:300:c641:1eff:fe3f:8cbb
0/ 100 = 0% |
2 4ms 0/ 100 = 0% 0/ 100 = 0% 2607:f500:0:1::a02:279
0/ 100 = 0% |
3 4ms 0/ 100 = 0% 0/ 100 = 0% 2607:f500:0:1::8e
0/ 100 = 0% |
4 14ms 0/ 100 = 0% 0/ 100 = 0% 2607:f500:0:1::8d
0/ 100 = 0% |
5 25ms 0/ 100 = 0% 0/ 100 = 0% et330.ash1.core.fuse.net [2607:f500:0:1::d844:e95]
Trace complete.
C:\Users\qdoct>

The top half of the output in Figure 12.33 looks similar to `tracert`. The bottom half is where you see the ping times and packet loss. In this example, there's no packet loss observed. Losing a small percentage of packets is normal and not a major concern; anything over 5 percent is considered high. Table 12.8 lists frequently used switches.

TABLE 12.8 `pathping` switches

Switch	Purpose
`-h` *number*	Defines the maximum number of hops to search. Useful if, for example, you just want to test connectivity to the ISP.
`-n`	Does not resolve each hostname. This speeds up the results.
`-p` *number*	Number of milliseconds to wait between pings. The default is 250. A good choice is 100. Again, speeds up results.
`-q` *number*	Number of queries per hop. The default is 100, choosing fewer speeds it up. Around 10–20 is usually enough.
`-w` *number*	Number of milliseconds to wait for each reply. The default is 3 seconds; 500 milliseconds is fine. Speeds up results.

Resolving Connectivity Issues

The whole purpose of using a network is to connect to other resources, right? So when networks don't work like they're supposed to, users tend to get a bit upset. The ubiquity of wireless networking has only made our jobs as technicians more complicated. In the following sections, we'll look at a variety of issues that you might run across and how to deal with them.

No Connectivity

Let's start with the most dire situation: no connectivity. Taking a step back to look at the big picture, think about all the components that go into networking. On the client side, you need a network card and drivers, operating system, protocol, and the right configuration. Then you have a cable of some sort or a wireless connection. At the other end is a switch or wireless router. That device connects to other devices, and so forth. The point is, if someone is complaining of no connectivity, there could be one of several different things causing it. As always, start with the basics and check the easy stuff first.

The most common issue that prevents network connectivity on a wired network is a bad or unplugged patch cable. Cleaning crews and the rollers on the bottoms of chairs are the most common threats to patch cables. In most cases, wall jacks are placed 4–10 feet away from the desktop. The patch cables are often lying exposed under the user's desk, and from time to time damage is done to the cable, or it's inadvertently snagged and unplugged. Tightly cinching the cable while tying it up out of the way is no better a solution. Slack must be left in the cable to allow for some amount of equipment movement and to avoid altering the electrical characteristics of the cable.

When you troubleshoot connectivity, start with the most rudimentary explanations first. Make sure that the patch cable is tightly plugged in, and then look at the card and check if any lights are on. If lights are on, use the NIC's documentation to help troubleshoot. More often than not, shutting down the machine, unplugging the patch and power cables for a moment, and then reattaching them and rebooting the PC will fix an unresponsive NIC.

A properly connected NIC should typically have one light illuminated (the link light). If the link light is not illuminated, it indicates a problem with the NIC, the patch cable, or the device to which the patch cable is connecting (hub, switch, server, and so on). Other lights that may be illuminated include a speed light, duplex light, and/or activity light.

If you don't have any lights, you don't have a connection. It could be that the cable is bad or that it's not plugged in on the other side, or it could also be a problem with the NIC or the connectivity device on the other side. Is this the only computer having problems? If everyone else in the same area is having the same problem, that points to a central issue.

Most wireless network cards also have indicators on them that can help you troubleshoot. For example, a wireless card might have a connection light and an activity light, much like

a wired network card. On one particular card we've used, the lights will alternate blinking if the card isn't attached to a network. Once it attaches, the connection light will be solid, and the link light will blink when it's busy. Other cards may operate in a slightly different manner, so be sure to consult your documentation.

If you don't have any lights, try reseating the wireless NIC. If you're using a USB wireless adapter, this is pretty easy. If it's inside your desktop, it will require a little surgery. If it's integrated into your laptop, you could have serious issues. Try rebooting first. If that doesn't help, see if you can use an expansion NIC and make that one light up.

No SSID Found

This is obviously a wireless problem, because wired networks don't use service-set identifiers (SSIDs). If a client can't find an SSID, it won't have network connectivity.

If a client can't find an SSID for a network it had previously joined, the router's SSID may have changed. If you're going to change an SSID, be sure to alert all users or technicians who might be affected. A second potential cause could be that the router is configured to not broadcast SSIDs. Some administrators will stop this broadcast, believing that doing so will increase the network security. It really doesn't, because any semi-decent hacker with a wireless packet sniffer will pick it up. What it does best, though, is to make it difficult for legitimate users to join the network.

A third potential cause is that the user's computer is too far out of range. Get closer to the access point or wireless router, and the signal should appear.

Let's assume that you have lights and that no one else is having a problem. (Yes, it's just you.) This means that the network hardware is probably okay, so it's time to check the configuration. Open a command prompt, type **`ipconfig`**, and press Enter. You should get an IP address. (If it starts with `169.254.` *`x.x`*, that's an APIPA address. We'll talk about those in the "Limited or Local Connectivity" section.) If you don't have a valid IP address, that's the problem.

The Linux, UNIX, and macOS version of `ipconfig` is `ip`.

Remember that in order to communicate on a network using TCP/IP (IPv4), you need to have a unique IP address and a valid subnet mask. If you want to communicate on a network outside of your own local network, you also need a default gateway.

If you do have a valid IP address, it's time to see how far your connectivity reaches. With your command prompt open, use the `ping` command to ping a known, remote working host. If that doesn't work, start working backward. Can you ping the outside port of your router? The inside port? A local host? (Some technicians recommend pinging your loopback address first with `ping 127.0.0.1` [or `ping ::1` on an IPv6 network] and then working your way out to see where the connectivity ends. Either way is fine. The advantage to starting with the loopback is that if it doesn't work, you know nothing else will either.) Using this methodology, you'll be able to figure out where your connectivity truly begins and ends.

The `ping` utility is useful for troubleshooting several types of connectivity problems that you might encounter. For example, if you can ping an IP address but not a hostname, then you know it's a DNS error. If you can ping a hostname (such as a remote web server) but you can't get to it with your web browser, then you know it's a problem with your browser or HTTPS or port 443, which may possibly involve the router configuration (like a blocked port).

Limited or Local Connectivity

In a way, limited connectivity problems are a bit of a blessing. You can immediately rule out client-side hardware issues because they can connect to some resources. You just need to figure out why they can't connect to others. This is most likely caused by one of two things: a configuration issue or a connectivity device (such as a router) problem.

Exam Objective 5.5, "Given a scenario, troubleshoot network issues," lists limited connectivity as a symptom to know.

Check the local configuration first. Use `ipconfig /all` to ensure that the computer's IP address, subnet mask, and default gateway are all configured properly. After that, use the `ping` utility to see the range of connectivity. In situations like this, it's also good to check with other users in the area. Are they having the same connectivity issues? If so, it's more likely to be a central problem rather than one with the client computer.

Real World Scenario

What Can (or Can't) You Get To?

A common refrain from users when they are having network issues is, "I can't get to anything." As a technician, you need to parse that statement to understand what they really can or can't get to, and if it's isolated to their device or occurs on other devices as well. It's best to ask questions to determine which resources are unavailable to them. The two main categories of unavailable resources are local resources and the Internet.

If the user is unable to get to local resources, odds are they can't get to "anything." Local resources include shared drives (aka shares), printers, or their email. In situations like this, treat it like there is no connectivity (or limited connectivity) and troubleshoot appropriately.

If local resources are accessible but the Internet is not, then you know the network card and wireless or wired connectivity is fine. Something is preventing the connection to the outside world. In these situations, there are a few specific things to check. First, check their IP address. If it starts with `169.254`, then they have an APIPA address, which means the DHCP server could be having issues. If they seem to have a legitimate IP address, can they ping an Internet website using its IP address? If so, but they can't ping the website name, then it's a DNS issue. If they can't ping the address, try a different one. Other devices that could potentially be blocking the traffic include routers, firewalls, and proxy servers.

APIPA and Link Local Addresses

As we talked about in Volume 1 Chapter 7, "Introduction to TCP/IP," Automatic Private IP Addressing (APIPA) is a service that autoconfigures your network card with an IP address. APIPA kicks in only if your computer is set to receive an IP address from the Dynamic Host Configuration Protocol (DHCP) server and that server doesn't respond. You can always tell an APIPA address because it will be in the format of `169.254.`*`x.x`*.

When you have an APIPA address, you will be able to communicate with other computers that also have an APIPA address but not with any other resources. The solution is to figure out why you're not getting an answer from the DHCP server and fix that problem.

Link local addresses are the IPv6 version of APIPA, and link local addresses always start with `fe80::` (they are in the `fe80::/10` range). They will work to communicate with computers on a local network, but they will not work through a router. If the only IP address that your computer has is a link local address, you're not going to communicate outside of your network. The resolution is the same as it is for APIPA.

IP Address Conflicts

Every host on a network needs to have a unique IP address. If two or more hosts have the same address, communication problems will occur. The good news is that nearly every operating system today will warn you if it detects an IP address conflict with your computer. The bad news is it won't fix it by itself.

The communication problems will vary. In some cases, the computer will seem nearly fine, with intermittent connectivity issues. In others, it will appear as if you have no connectivity.

The most common cause of this is if someone configures a computer with a static IP address that's part of the DHCP server's range. The DHCP server, not knowing that the address has been statically assigned somewhere, doles out the address and now there's a conflict. Rebooting the computer won't help, nor will releasing the address and getting a new lease from the DHCP server—it's just going to hand out the same address again because it doesn't know that there's a problem.

As the administrator, you need to track down the offending user. A common way to do this is to use a packet sniffer to look at network traffic and determine the computer name or MAC address associated with the IP address in question. Most administrators don't keep network maps of MAC addresses, but everyone should have a network map with hostnames. If not, it could be a long, tedious process to check everyone's computer to find the culprit.

Usually the person who manually configured their address didn't intend to cause any problems. This would be a good time to show your professionalism and communication skills and educate the user as to why they shouldn't have done what they did.

Intermittent Connectivity

Intermittent connectivity is when the network sometimes connects, but it's not consistently connected. Sometimes the connection will quickly disappear and reappear, and other times it will be disconnected for longer—minutes can seem like hours when this happens. This category of problems is the most difficult and frustrating to troubleshoot. Getting to why it sometimes behaves but other times does not isn't always easy. Under this heading, we're going to consider intermittent connectivity (wireless and Internet), slow network speeds, high latency, and related issues, because they are all pretty similar. Terms like "slow network speeds" are fairly straightforward, but not all terms in this section are. Let's define a few, and then get into more detail, because most of these problems have the same causes and resolutions:

- *Latency* is a measure of delay on a network. High latency looks like the connection is intermittent, simply because the delays are so long. If latency is happening on a large data transfer, you might not notice unless the connection terminates. If it happens during video streaming, or on a virtual meeting or a *Voice over Internet Protocol (VoIP)* phone call, you will definitely notice because the person or people on the other end of the line will freeze or completely cut out.
- *Jitter* is variable latency. Think of it as if someone had too much caffeine and they are shaking a bit. Sometimes they are able to hold still, other times they are . . . jittery. Same thing here.
- *Port flapping* is when a switch port opens and closes very quickly. According to Cisco, it's technically when the port opens and closes three or more times in a second for at least 10 seconds. The result is intermittent connectivity.

On a wired network, if you run into slow speeds or intermittent connectivity, it's likely a load issue. There's too much traffic for the network to handle, and the network is bogging down. Solutions include adding a switch, replacing hubs with switches (if your network is using hubs in today's age, it's *really* time to upgrade), and even creating virtual LANs (VLANs) with switches. If the network infrastructure is old (for example, if it's running on Category 5 cable or you only have 100 Mbps switches), then it's definitely time for an upgrade. Remember that your network can only be as fast as its slowest component.

Exam Objective 5.5 has several symptoms that are covered in this section. They include intermittent wireless connectivity, slow network speeds, jitter, poor VoIP quality, port flapping, high latency, external interference, and intermittent Internet connectivity.

Other wired issues can include bad or poorly connected cables or faulty switch ports. Check to ensure that the cable is properly connected on both ends. Most cables today have a latch that holds them into place, but it's still a good idea to detach and reattach them for good measure. You can always try a different cable, as well, to see if that resolves the problem. If the cable seems to be fine, then perhaps try a different port on the switch or hub, if possible. Finally—and this is rare in today's networking—there could be a speed or duplex mismatch between the sending and receiving devices. Network adapters and connectivity devices can be set to limit the speeds, as a form of backward compatibility. If one device is set too slow, then another transmitting at full speed may have issues connecting with it. With duplex, most NICs are fully capable of sending and receiving at the same time. That's called full duplex. Some older cards were able to operate at only half duplex, meaning they could send or receive, but not both at the same time. Today, nearly all devices autodetect speed and duplex, so this is, as we said, a rare issue. If all else fails, though, it never hurts to check the configuration settings.

Again, the previous steps will fix most issues, but there a few more things to know about VoIP and port flapping:

Fixing Poor VoIP Quality Most corporate landline telephones today run over an IP network. This is called *Voice over Internet Protocol (VoIP)*. A bad or intermittent network connection will definitely cause poor VoIP quality. Most routers have an option called Quality of Service (QoS) that allows administrators to set priorities for different streams of network traffic. If a network is suffering from poor VoIP quality, the administrator could give it a higher priority. That way, if the network connection is overloaded, higher-priority traffic should be less affected.

Dealing with Port Flapping If confronted with port flapping, the advice given earlier still applies—check the cables, connections, and look for a speed or duplex mismatch. There are two other things to check for as well. First, look at the error logs for the switch. There should be a layer 1 issue on both ends of the link. (Remember the OSI model?) If not, it could be an intermediary device between the ends of the link causing the problem. Second, reseat the small form-factor pluggable (SFP) connectors (one is shown in Figure 12.34). SFPs are pinky-sized, hot-swappable transceivers used in gigabit and faster connections. If reseating doesn't help, switch out the SFPs on the link with ones that are known to work.

Wireless networks can get overloaded, too. It's recommended that no more than 30 or so client computers use one wireless access point (WAP) or wireless router. (Wi-Fi 6 and newer can handle more but should still be limited to 60–100 at most.) Any more than that can cause intermittent access problems. The most common reason that users on wireless networks experience any of these issues, though, is distance. The farther away from the WAP the user gets, the weaker the signal becomes. When the signal weakens, the transfer rates drop dramatically.

For example, the signal from an 802.11ac (Wi-Fi 5) wireless router has a maximum indoor range of about 35 meters (115 feet), barring any obstructions. At that distance, though, 802.11ac will support transfer rates of only about 50 Mbps—far less than the 1.3 Gbps the users think they're getting. External interference, such as from radio signals, microwaves, large motors, and fluorescent lights, as well as physical barriers such as concrete or steel, can greatly reduce the effective range. The solution is to move closer or install more access points. Depending on the configuration of your working environment, you could also consider adding a directional antenna to the WAP. The antenna will increase the distance the signal travels but only in a limited direction.

FIGURE 12.34 Gigabit SFP

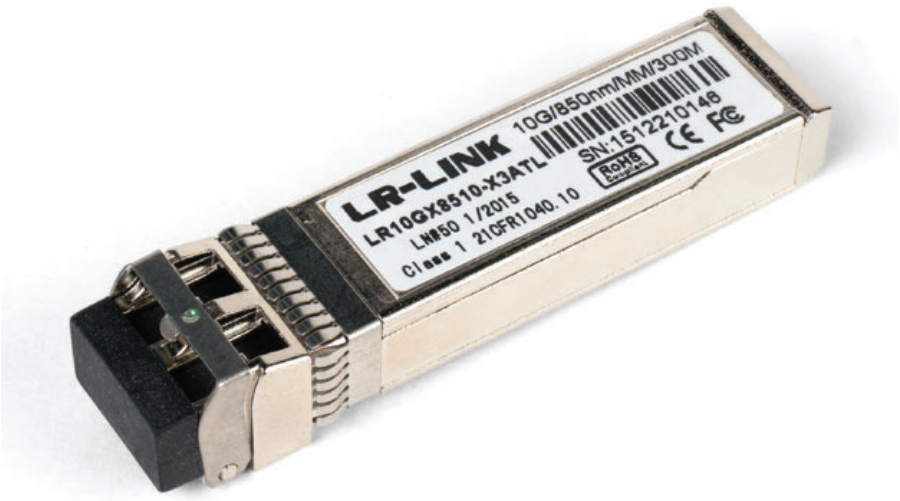

Source: Dmitry Nosachev / Wikimedia Commons / CC BY SA 4.0

On a wireless network, intermittent or no connectivity issues can sometimes be resolved by disconnecting from (and forgetting) the network and then re-joining. Turning the wireless adapter off and back on again might help as well.

Troubleshooting intermittent Internet connectivity should be handled the same way as an internal connection. The only difference is that there are more devices involved, including the router, maybe a firewall, other routers on the Internet, and the system hosting the website the user is trying to access. See if other websites work better or faster

Authentication Failures

If a user can't authenticate on the network, they will not be able to access any resources. Authentication in a Windows Server environment is handled by Active Directory domain controllers.

Authentication failures are listed as a symptom to know in Exam Objective 5.5.

The most common cause of authentication failures is users typing in the incorrect password. Check to make sure the Caps Lock key is not on. If it's not user error, then check to ensure that the authentication server is online and reachable. Is this user the only one having problems authenticating, or are other users affected as well? Pinging the server from the affected computer or another one nearby may help answer the question.

Summary

This chapter kicked off our discussion of the troubleshooting objectives for the A+ 220-1201 exam. We briefly touched on some troubleshooting tricks and tips as well as provided a systemic framework for troubleshooting. Both art and science are involved, and experience in troubleshooting is helpful but not a prerequisite to being a good troubleshooter. You learned that in troubleshooting, the first objective is to identify the problem by narrowing the possibilities to isolate the cause.

Next, we investigated the causes and symptoms of core hardware problems, such as unusual noise, excessive heat, burning smells and smoke, visible damage, and intermittent device failure. After the discussion of general hardware, we talked about issues specific to internal components, including the motherboard, CPU, RAM, and power supply.

After that, we covered troubleshooting storage drives and RAID arrays. We talked about lights, sounds, and audible alarms, devices not being found, performance issues, S.M.A.R.T. diagnostics, RAID issues, and optical drive issues.

The chapter finished with network troubleshooting. You learned about various hardware and software tools you can use to troubleshoot networking issues. Then, you learned about resolving connectivity issues as well as authentication failures.

Exam Essentials

Understand what happens during the POST routine. During the power-on self-test (POST), the BIOS/UEFI checks to ensure that the base hardware is installed and working. Generally, one POST beep is good. Any more than that and you might have an error.

Understand problems related to the system BIOS/UEFI. BIOS/UEFI settings are maintained by the CMOS battery when the system is powered off. If the system keeps losing the date and time or boot settings, it could indicate a problem with the CMOS battery.

Know what is likely to cause unexpected shutdowns, system lockups, continuous reboots, and intermittent device failures. All these issues can be caused by a failing motherboard, CPU, or RAM. In the case of other intermittent device failures, it could be that specific device as well. Many times these issues are exacerbated by overheating.

Understand common problems that power supplies can cause. Power supplies can fry components, but they can also cause no power, grinding or squealing noises, spinning fans but no power to other devices, smoke, and burning smells.

Know which devices within a system can make unusual noises. Unusual noises are not pleasant to hear. Generally speaking, only devices with moving parts, such as HDDs, power supplies, and fans, can produce unwanted unusual noises.

Know how to avoid overheating. Using fans and heat sinks will help to avoid overheating. Also know that overclocking the processor can cause overheating.

Know the proprietary crash screens for Windows and Mac operating systems. Windows has the Blue Screen of Death (BSOD), whereas macOS uses the pinwheel.

Understand what a distended capacitor is. It's when a capacitor swells and possibly bursts, releasing a reddish-brown electrolyte. A motherboard with distended capacitors will likely fail. Don't touch the electrolyte!

Know what causes a black screen. Black screens are likely the fault of the video card, video cable, or display unit. If the video circuitry is built into the motherboard, it could be a faulty motherboard as well.

Understand causes of sluggish performance. Generally, sluggish performance is related to the memory or hard drive. If either are being overused, the system will be slower to respond. Overworked CPUs can also cause sluggish performance.

Know what causes application crashes. App crashes are most likely one of two things: a poorly coded app or faulty memory.

Understand common issues related to hard drives and RAID arrays. There are several issues that are hard drive–related. The drive can have a read/write failure or slow performance. If a conventional drive fails, it might produce a loud clicking or scratching noise. Sometimes boot failures are hard drive–related, if the drive can't be found or recognized or the OS isn't found due to a bad boot sector or master boot record. If the OS crashes, it may produce a proprietary crash screen, such as the BSOD in Windows or the pinwheel in macOS. RAID arrays are collections of hard drives and can have similar failures and issues with not being found. All hard drives made today have S.M.A.R.T. diagnostics built into them.

Know what to check if your wireless networking card isn't working. Make sure the card has lights indicating that it's working. You might also have an external toggle switch to turn the card on and off. Finally, if your computer has an external RJ-45 connection, you can plug it in and see if it works when wired.

Understand what various networking hardware tools are used for. Crimpers are for cutting and stripping cables and putting the connectors on the ends. Cable strippers are similar but don't put connectors on. A Wi-Fi analyzer examines wireless network signals. Toner probes allow you to trace a cable from one point to another. Punch-down tools connect wires to frames, such as a 110 block. A cable tester allows you to verify that a cable works; loopback plugs test the functionality of network cards. A network tap creates a copy of network traffic for use by monitoring devices without disrupting normal traffic.

Know what the `ipconfig`, `ping`, and `tracert` commands are used for. Admittedly, these are specifically for A+ exam 220-1202, but know what they do. Both `ipconfig` and `ping` are network troubleshooting commands. You can use `ipconfig` to view your computer's IP configuration and `ping` to test connectivity between two network hosts. (The `ip` command in macOS and Linux is equivalent to `ipconfig` in Windows.) `tracert` allows you to view the network path a packet takes from the host to the destination.

Know what the `netstat`, `net`, and `nslookup` commands are used for. These commands need to be understood for Exam 220-1202 as well. `netstat` shows network statistics; `net` allows you to perform network-management tasks, such as sharing folders; and `nslookup` allows you to query a DNS server.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. When a user turns on their desktop computer, nothing appears on the screen. They hear three long beeps, followed by three short beeps, and three long beeps. What should you do as a next step?
   - **A.** Replace the motherboard.
   - **B.** Replace the RAM.
   - **C.** Replace the CPU.
   - **D.** Look up the beep code on the manufacturer's website.
2. A field technician reports back that the computer they are troubleshooting has exposed distended capacitors. Which component has most likely failed?
   - **A.** Power supply
   - **B.** Motherboard
   - **C.** CPU
   - **D.** RAM
3. You are troubleshooting a Windows-based desktop computer that is experiencing unexpected shutdowns. Which of the following would be a good next step to resolve the issue? (Choose two.)
   - **A.** Replace the motherboard, CPU, and RAM to see if it resolves the problem.
   - **B.** Replace the RAM to see if it resolves the problem.
   - **C.** Reinstall Windows to see if it resolves the problem.
   - **D.** Replace the CPU to see if it resolves the problem.
4. When a desktop computer boots up, there is no video on the LED monitor, just a blank screen. The system makes one beep when it powers on, and air is coming out of the case fan. What is the most likely cause of the problem?
   - **A.** The motherboard has failed.
   - **B.** The video card has failed.
   - **C.** The monitor is disconnected.
   - **D.** The CPU or RAM has failed.
5. You are training a class of new technicians on troubleshooting motherboard issues. Which of the following are issues most likely caused by a failing CMOS battery? (Choose two.)
   - **A.** BIOS time and setting resets
   - **B.** Attempts to boot to the incorrect device
   - **C.** Intermittent device failure
   - **D.** BSOD/pinwheel errors

6. A desktop computer is experiencing continuous reboots. It gets to the BIOS screen and then reboots every time. Which two components are most likely to cause this problem? (Choose two.)
   - **A.** Power supply
   - **B.** RAM
   - **C.** Motherboard
   - **D.** CPU

7. A user pushes the power button on their desktop computer but nothing happens. They double-check that it is plugged in and that the wall outlet is working properly. Which components are most likely to cause this problem? (Choose two.)
   - **A.** Power supply
   - **B.** Motherboard
   - **C.** Power cord
   - **D.** CPU
   - **E.** RAM

8. A technician who has been troubleshooting a computer tells you that the system is having problems with overheating. Which component is most likely to cause this issue?
   - **A.** Power supply
   - **B.** Motherboard
   - **C.** CPU
   - **D.** RAM

9. You turn on a desktop computer and it fails to boot. You hear a rhythmic ticking coming from inside the case. What is most likely the problem?
   - **A.** Failed hard drive
   - **B.** Failed motherboard
   - **C.** Failed CPU
   - **D.** Failed RAM
   - **E.** A bomb

10. When you turn on a desktop computer, you hear a loud squealing noise coming from inside the case. Which components are potential culprits for making the noise? (Choose all that apply.)
    - **A.** CPU fan
    - **B.** Power supply fan
    - **C.** HDD
    - **D.** SSD
    - **E.** CPU
    - **F.** RAM

11. A user complains of smoke and a foul odor coming from their desktop computer, so they shut it down. Which component is most likely causing the problem?

A. Power supply

B. RAM

C. CPU

D. Motherboard

12. A Windows user complains that their computer keeps crashing. What is the proprietary crash screen they are most likely seeing?

A. Stop sign

B. Pinwheel

C. BSOD

D. Black screen

13. A laptop user with an 802.11ax wireless network card just switched desk locations. In the new location, they say that their wireless drops connectivity intermittently and seems to have slow transfer speeds. Which of the following is most likely to resolve these issues?

A. Replace the network card.

B. Run `ipconfig` to ensure that they are receiving the right IP address from the DHCP server.

C. Install a new WAP closer to their new desk.

D. Use a network monitor to ensure that there are no IP address conflicts.

14. While troubleshooting a client computer, you decide to obtain a new IP address from the DHCP server. After releasing the existing address, which command do you use to get new IP information from the DHCP server?

A. `ipconfig /refresh`

B. `ipconfig /renew`

C. `ip /release`

D. `ip /start`

15. You believe that a network card is beginning to fail. Which of the following tools should you use to see if the network card is capable of sending and receiving data from an RJ-45 port?

A. Toner probe

B. Loopback plug

C. Multimeter

D. Cable tester

16. Troubleshooting a network connectivity issue takes you into the wiring closet, where you realize that none of the cables are properly labeled. Several are disconnected. Which tool should you use to determine where the cables run to?
    - **A.** Loopback plug
    - **B.** Punch-down tool
    - **C.** Cable tester
    - **D.** Toner probe

17. You turn on a laptop computer and the hard drive is not recognized. Where should you go to troubleshoot the issue?
    - **A.** S.M.A.R.T. diagnostics
    - **B.** BIOS/UEFI
    - **C.** Windows Device Manager
    - **D.** Boot sector manager

18. A network cable for a crucial server has failed. There are no pre-made Ethernet cables available, but there is bulk cable and connectors. Which tool do you need to use to make a new cable?
    - **A.** Crimper
    - **B.** Punch-down tool
    - **C.** Cable tester
    - **D.** Loopback plug

19. You are troubleshooting a server and discover that one of the hard drives in the RAID 0 array has failed. Which statement is true?
    - **A.** You need to replace the failed drive, but the data is okay because the drive is configured as a mirror.
    - **B.** You need to replace the failed drive, but the data is okay because the drive is configured as a disk stripe with parity.
    - **C.** You need to replace the failed drive, and the data on the array is lost.
    - **D.** You do not need to replace the failed drive; the system will function normally.

20. While troubleshooting an intermittent Internet connectivity problem, you notice that port 443 on the router keeps opening and closing. What is this an example of?
    - **A.** Jitter
    - **B.** High latency
    - **C.** External interference
    - **D.** Port flapping

Performance-Based Question

You will encounter performance-based questions on the A+ exam. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answers compare to the authors', refer to Appendix B.

Match the networking tools to the jobs they perform.

TOOLS

JOBS

1. Put connectors on UTP.
2. Connect wires to a 110-block.
3. Cut cables.
4. Monitor network traffic.
5. Make sure a new cable works.
6. Test if a NIC is working.
7. Trace a cable from end to end.
8. Determine the strength of Wi-Fi signals.

Chapter 13

Mobile Device, Display, and Printer Troubleshooting

THE FOLLOWING COMPTIA A+ 220-1201 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

✓ **5.0 Hardware and Network Troubleshooting**

✓ **5.3 Given a scenario, troubleshoot video, projector, and display issues.**

- Common symptoms
  - Incorrect data source
  - Physical cabling issues
  - Burnt-out bulb
  - Fuzzy image
  - Display burn-in
  - Dead pixels
  - Flashing screen
  - Incorrect color display
  - Audio issues
  - Dim image
  - Intermittent projector shutdown
  - Sizing issues
  - Distorted image

✓ **5.4 Given a scenario, troubleshoot common mobile device issues.**

- Common symptoms
  - Poor battery health

- Swollen battery
- Broken screen
- Improper charging
- Poor/no connectivity
- Liquid damage
- Overheating
- Digitizer issues
- Physically damaged ports
- Malware
- Cursor drift/touch calibration
- Unable to install new applications
- Stylus does not work
- Degraded performance

✓ **5.6 Given a scenario, troubleshoot printer issues.**

- Lines down the printed page
- Garbled print
- Paper jams
- Faded prints
- Paper not feeding
- Multipage misfeed
- Multiple prints pending in queue
- Speckling on printed pages
- Double/echo images on print
- Grinding noise
- Finishing issues
  - Staple jams
  - Hole punch
- Incorrect page orientation
- Tray not recognized
- Connectivity issues
- Frozen print queue

Hardware problems are sometimes very easy to identify. If you push the power button and nothing happens, you can be pretty confident that it's not the fault of the operating system. Other hardware problems are more complicated. For example, memory issues may cause errors that look like they're the fault of an application, or what appears to be a failing device may be resolved by updating the software driver. So while this chapter focuses on troubleshooting hardware, understand that in many cases you'll need to know how to navigate through apps and software settings to really narrow down the issue.

With all of the integration between software applications and hardware components, it can be challenging to understand where one stops and the other starts, or how their interoperation affects one another. To top it all off, you're probably going to be working in an environment that requires you to understand not just one computer but a network full of workstations, servers, switches, routers, and other devices, and how they should play nicely together.

This introduction isn't intended to scare you, but rather to point out that computers are complicated and troubleshooting can be hard. You already know that, but give yourself some latitude in troubleshooting, because it's not always easy. Prepare yourself with a solid foundation of knowledge and never be afraid to consult other resources, such as the Internet or coworkers. Situations will arise that make even the most experienced technicians shake their heads in frustration.

As you learned in Volume 1 Chapter 12, "Core Hardware and Network Troubleshooting," the best way to tackle any problem is to take a systematic approach to resolving it. This applies to the mobile device, display, and printer issues that we'll talk about here as well as the software and security issues that we will cover in Volume 2 Chapter 7, "Troubleshooting Operating Systems and Security."

Recall from Volume 1 Chapter 12 that while experience helps, novice technicians can be effective troubleshooters as well. Troubleshooting becomes a lot easier if you follow logical procedures to help you develop experience. The first thing to do is always to check the easy stuff, such as physical cables and connections. You would be amazed at how many times the simple question "Is it plugged in?" resolves hardware problems. Second, see if anything has recently changed or if there are any recent incidents that might have caused the problem. For example, if someone's laptop won't boot up, you might not have a clue as to why. But if they tell you that they just dropped it down the stairs, you might have a better idea of where to start. Finally, narrow down the scope of the problem. Find out exactly what works and what doesn't. Knowing where the problem starts and stops helps you to focus your troubleshooting efforts.

This chapter finishes the hardware troubleshooting discussion we started in Volume 1 Chapter 12. Specifically, we will look at the following:

- Video systems, with an emphasis on projectors and displays
- Mobile devices, including laptops
- Printers

You may wonder why we broke hardware troubleshooting into multiple chapters. If we're being honest, it's due to the large number of problems that could arise, the volume of A+ exam objectives that creates, and the amount of room it takes to discuss them all. In other words, we didn't want to hit you with a massively long chapter!

Even with breaking it up, though, there's no way that a reasonably sized book could teach you about all the possible issues you could face. Nor would it be logical for you to try to memorize them all. Instead, we focus on some of the common issues and help you think through the process of narrowing down the possibilities so that you can efficiently resolve any problem you encounter. Armed with that knowledge, you should have the confidence to tackle any problem you face, whether or not it's familiar to you.

Troubleshooting Video, Projector, and Display Issues

Troubleshooting video problems is usually fairly straightforward because there are only a few components that could be causing the problem. You can sum up nearly all video problems with three simple statements:

- There is either no video or bad video.
- It could be user error, but if not, then . . .
- The video card or the monitor/projector/display/cable is to blame.

In the majority of cases when you have a video problem on a desktop computer, a good troubleshooting step is to check the monitor by transferring it to another machine that you know is working. See if it works there. If the problem persists, you know it's the monitor. If it goes away, you know it's the video card (or possibly the driver). Is the video card seated properly? Is the newest driver installed?

CompTIA A+ 220-1201 Exam Objective 5.3, "Given a scenario, troubleshoot video, projector, and display issues," lists 13 symptoms you should understand and know how to fix. We'll break them into three categories:

- Input issues
- Image problems
- Other issues

Let's take a look at each of them now.

Video Input Issues

Imagine you're getting ready for a big presentation. Everyone is gathered in the room, and you connect your laptop to the video projector or external monitors—and there's no display. The audience sighs and people start getting fidgety or multitasking. It's not a great situation. Odds are that if this hasn't happened to you, then you've been on the opposite side—sitting there wishing the presenter could figure it out. What's the best way to resolve this type of issue? Try these three steps:

1. Check the obvious. (Where have you heard this before?) Is everything plugged in and powered on? Look for indicator lights on the projector or external monitor to be sure.
2. Make sure the laptop is set to output video to the right device. On most laptops, you need to press the function key and another key known as the *LCD cutoff* switch (such as F4 or F8) to direct the video output to an external monitor. This is called *toggling the display*. You might need to do this a few times. Figure 13.1 shows portions of two laptop keyboards—in the top one, F7 is the toggle, and in the bottom one, it's F8. (Later, in Figure 13.6, you will see a monitor icon on F4 that does the same thing.)

Step 2 in checking video input issues mentions the LCD cutoff switch. You will hear this called a variety of names, such as an *LCD toggle* or *display toggle*. Note that it's named LCD because the original laptop displays were LCDs, and the legacy name carries over to today. In this chapter, you will see us use the term LCD and in general we're referring to all modern, flat-screen displays such as LEDs and OLEDs too. We will call out specific differences between flat-screen displays when it's needed.

3. Verify that the monitor or projector is configured for the correct data input source. Having an incorrect data source means that the device won't display anything. Most monitors and projectors have multiple inputs, such as three or four HDMI, DisplayPort, and possibly even VGA ports. There may be a button on the monitor or projector to bring up an onscreen menu and configure the correct source, or you may need to use the device's remote control.

Incorrect input source is the first of the symptoms you need to know as part of Exam Objective 5.3. Also remember that a "no display" issue can be caused by physical cabling issues (another symptom on the list).

If everything checks out, it's possible there could be physical cabling issues. You can try disconnecting and reconnecting the video cables or another cable if possible.

FIGURE 13.1 LCD cutoff switches (video toggle keys), F7 on the top keyboard and F8 on the bottom keyboard

Video Image Problems

This group of symptoms deals specifically with the image on the screen, or the lack thereof. Here are the ones you should know:

Burnt-Out Bulb This is a projector-specific issue. Projector bulbs have a limited shelf life, and the heat produced by a projector and its bulb can shorten it even further. If the bulb burns out, then clearly there will be nothing displayed. The only resolution is to replace the bulb. Be sure it cools down first before you try to replace it, though—as mentioned, they run hot and they will burn you.

All of the symptoms in the "Video Image Problems" section are listed among the symptoms in Exam Objective 5.3. You should know how to troubleshoot the following:

- Burnt-out bulb
- Fuzzy image
- Sizing issues
- Distorted image
- Display burn-in
- Dead pixels
- Flashing screen
- Incorrect color display
- Dim image

Fuzzy Image Resolving a fuzzy image problem will differ depending on the display device. For example, projectors have focus mechanisms that allow them to produce images on screens at different distances. A lot of projectors will try to autofocus but will

also have onscreen menu options or a knob around the outside of the lens to manually adjust the focus.

A fuzzy image on an LCD screen is an entirely different story. It could be caused by external interference such as fluorescent lights, magnetic devices, and electrical devices such as fans, lamps, and speakers. Check for any of those nearby. If the display uses a cable, it could also be a loose or bad cable. Finally, it could be that the resolution is set for something that the display can't handle, or at least can't handle well. (To be fair, most of the time if the resolution isn't supported, the image will appear warped or stretched and distorted, but it could be fuzzy as well.)

There are a few things you can try to fix it. In Windows, right-click an open area of the desktop and choose Display Settings (Figure 13.2). Once there, you can change the display resolution. Another thing to try is to click Scale to get to Custom Scaling (Figure 13.3). Custom scaling features, configured on this page, could cause fuzziness as well.

If none of those things work, check the manufacturer's website for similar issues, but it could be that the display itself is bad.

FIGURE 13.2 Windows Display settings

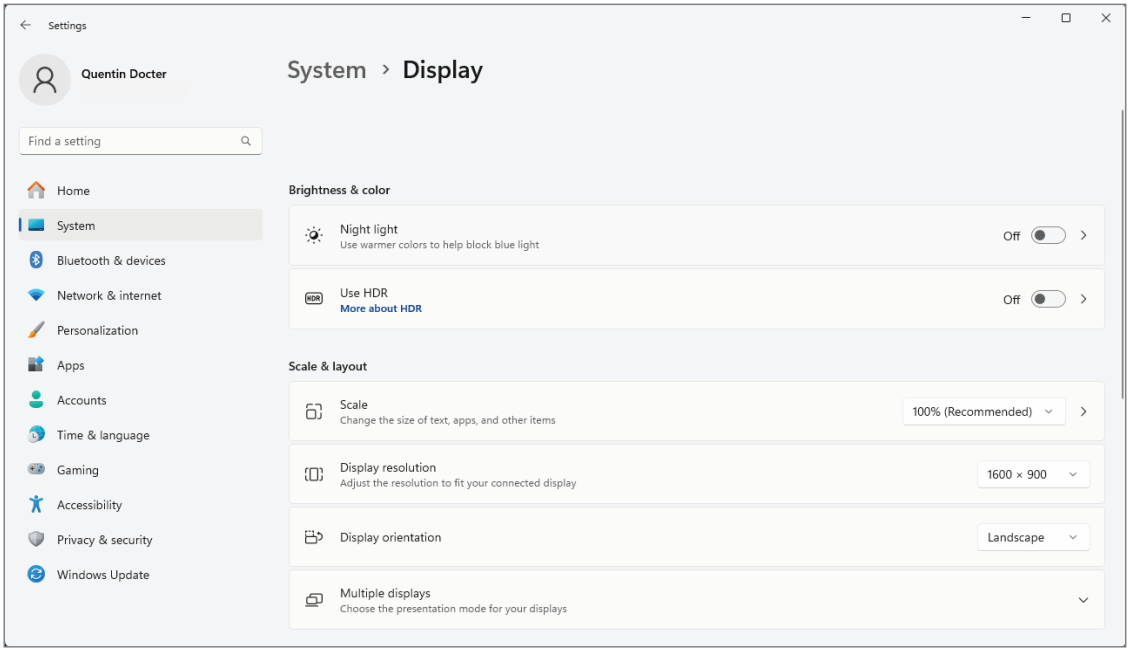

Sizing Issues Since we just discussed Windows Display settings, now is a good time to point out that sizing issues can be affected by the screen resolution as well. If the resolution is set too high, the icons and text on the screen could be too small to read. The

opposite can happen too, where the icons seem enormous. In that case, the resolution might be too low and could be increased. Right-click on the desktop and go into Display settings to fix this.

FIGURE 13.3 Scaling settings

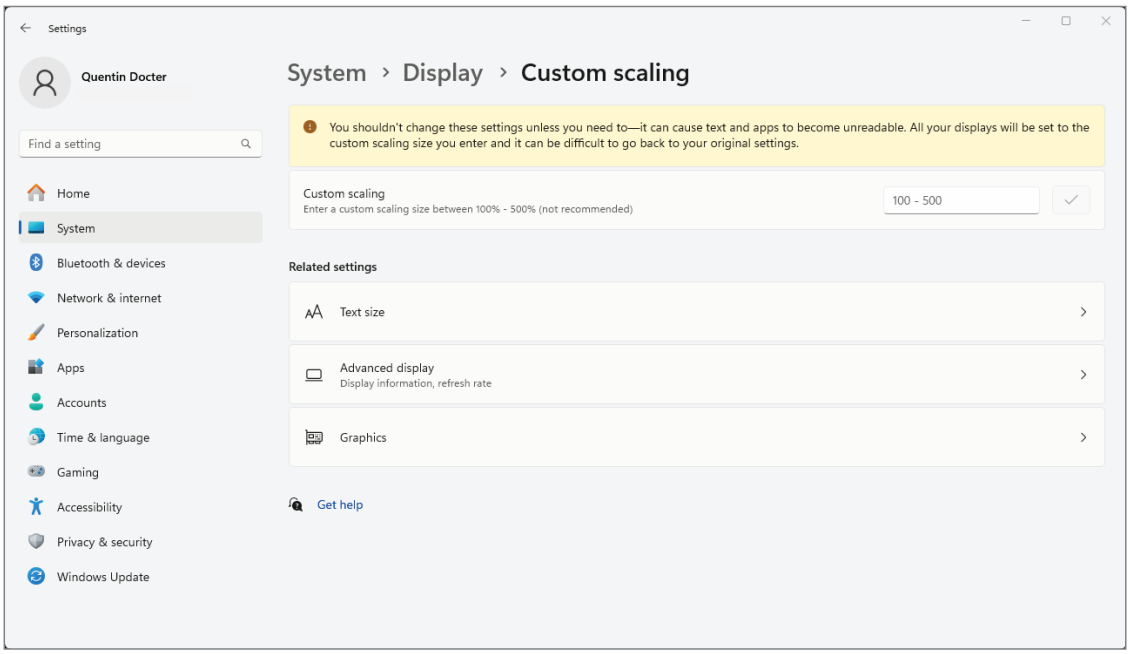
Settings
Quentin Docter
Find a setting
Home
System
Bluetooth & devices
Network & internet
Personalization
Apps
Accounts
Time & language
Gaming
Accessibility
Privacy & security
Windows Update
System › Display › Custom scaling
You shouldn't change these settings unless you need to—it can cause text and apps to become unreadable. All your displays will be set to the custom scaling size you enter and it can be difficult to go back to your original settings.
Custom scaling
Enter a custom scaling size between 100% - 500% (not recommended)
100 - 500
Related settings
Text size
Advanced display
Display information, refresh rate
Graphics
Get help

Distorted Image This is the third symptom in a row that could be resolution-related. Recall from Volume 1 Chapter 3, "Peripherals, Cables, and Connectors," that most displays have a preferred aspect ratio they operate under, such as 16:9 or 16:10. The default resolution of almost every digital display will be one of these two aspect ratios. If a display wants to be set at 16:9 resolution (which is most common) but it's configured for a 16:10 or 4:3 resolution, the image on the screen could appear distorted.

To check this, go into Display settings and look at the default resolution (you can see it on the menu of resolutions in Figure 13.4). The recommended resolution is 2560 × 1440, which is 16:9. (You can tell by dividing the first number by the second. Both 2560 ÷ 1440 and 16 ÷ 9 equal 1.78.) If you were to set the display to another ratio, such as 1600 × 1024, one of two things will happen. Either the monitor will stretch or compress the image to display it on the entire screen—causing a distorted image—or the image on the display will have blank borders along the top and bottom or the sides as it tries to maximize the size of the image without distortion. Either way, the solution is to change the resolution back to something the display natively supports.

FIGURE 13.4 Display resolutions

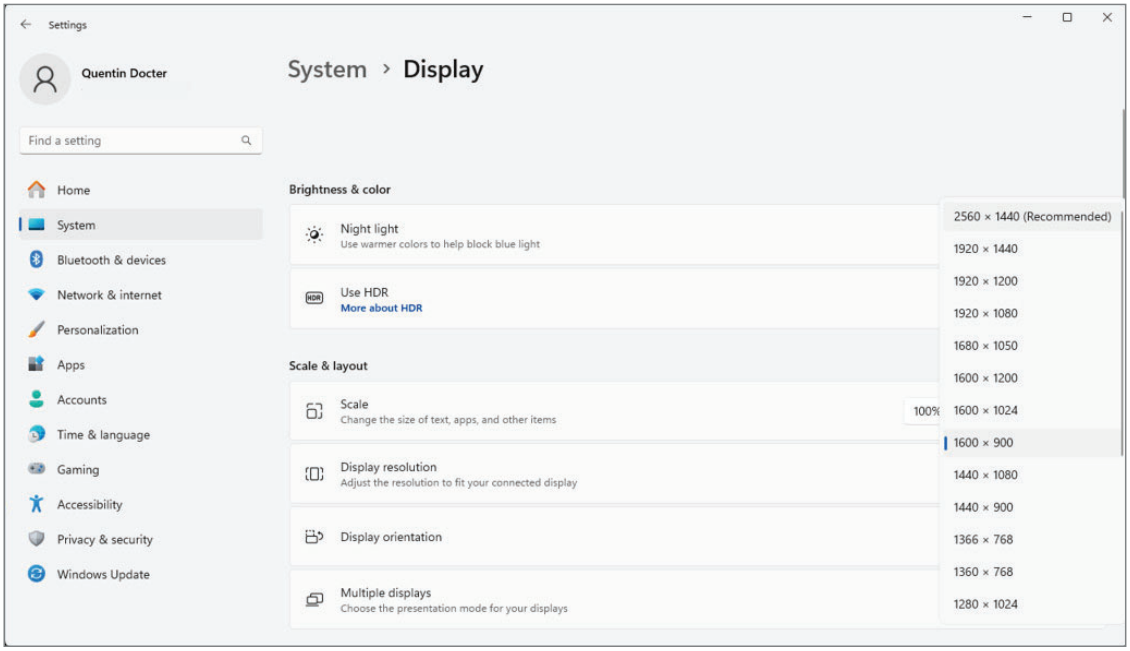

Some distorted image issues are not related to resolution. For example, if the display image is oddly shaped, like a trapezoid or hourglass, the display device could be going bad. Try turning it off and back on again or opening its onscreen menu (using a button on the bottom, side, or back of the unit) to see if there's a factory reset or adjustment you can make. If nothing seems to work, it could be a bad display unit.

Display Burn-In When this happens, no matter what you have on your screen, you can still see the outlines of a different image. That image has been "burned" into the monitor (sometimes simply referred to as *burn-in*) and isn't going away. Sometimes you will also hear this referred to as an artifact. Regardless, the only solution is to replace the monitor.

Older LCD and plasma monitors were susceptible to burn-in, but most modern display technologies aren't. The exception is OLED displays, because the individual organic pixels can degrade over time if static images are displayed for extended periods of time.

Dead Pixels Dead pixels are spots on the screen that never "fire," or light up. You can check for these by setting the background to white and seeing if any spots don't light up. The only solution for dead pixels is to replace the monitor.

Dead pixels never light up, but the opposite can happen too, where a pixel is "stuck" in an on position. So instead of it being a black dot on a light screen, it's a lit dot on a black screen. Some suggest that to fix it, you can try to toggle the colors on and off quickly (black screen to white

and vice versa), or you can press on the display in that area to get it "unstuck." Don't press too hard—a gentle massage should do. Personally, we've never had success with either of these two fixes, but they could be worth a shot. Another option is a utility such as JScreenFix (`www.jscreenfix.com`), which may work as well but isn't guaranteed. If a pixel won't unstick, then the solution is to replace the monitor.

Flashing Screen Sometimes a display will either subtly flicker or flash off and on. Those two symptoms are caused by different things. Flickering screens could be caused by having the screen refresh rate set too low. Recall from Volume 1 Chapter 3 that the refresh rate is essentially how often the screen is re-drawn. If set too low (usually 60 Hz or lower), it can cause the screen to flicker or appear to roll. To change the refresh rate, go to Display Settings ➤ Advanced Display, as shown in Figure 13.5.

FIGURE 13.5 Advanced Display refresh rate

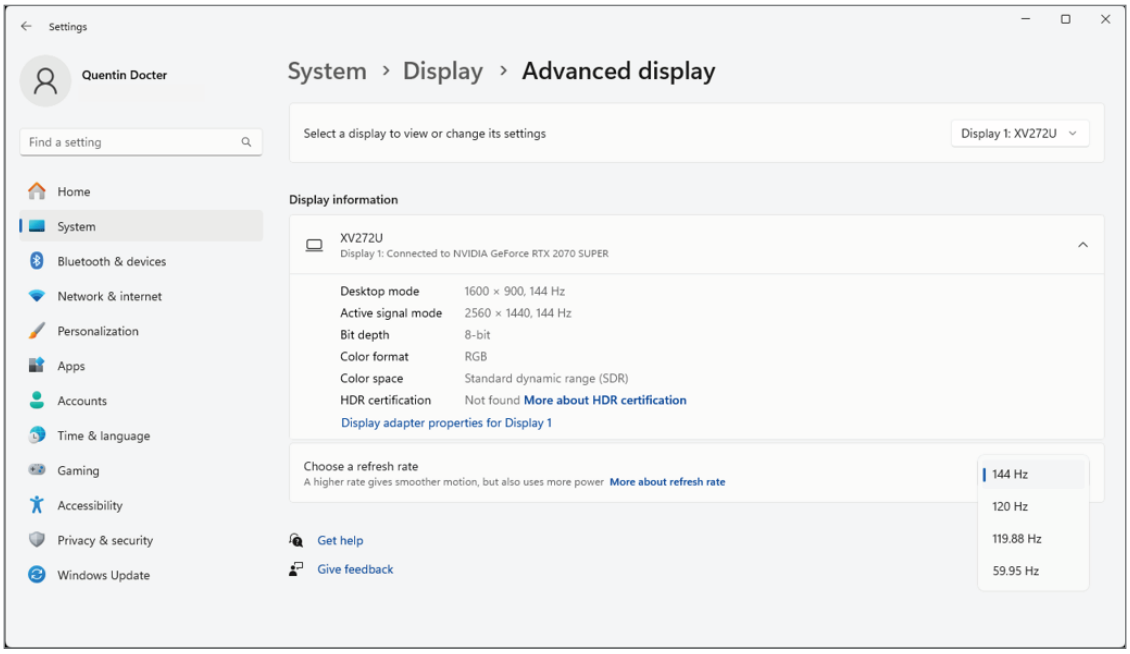

From a hardware standpoint, flickering screens are most commonly caused by the backlight starting to fail. Also, recall that some displays use an inverter to provide power to the backlight. Backlights are more prone to failure than inverters, but either one could be the problem. If either one of those are failing, replace the backlight or the inverter.

Flashing off and on could be the backlight, but it could also be a loose cable or an unsupported resolution. Try the usual fixes, including checking the cables (if applicable), changing the resolution, or reinstalling the video card driver.

Incorrect Color Display This can happen when the LCD monitor's controller board starts to fail and doesn't perform color mapping correctly. It also happens on old CRTs. To fix it on an old CRT display, you use a process called *degaussing* (decreasing or eliminating an unwanted magnetic field), which is done through a utility built into the menu on the monitor. Finally, this can also happen if the pins on the connector are damaged or if the connector isn't plugged in all the way. If switching the monitor makes the problems go away, it's probably time to replace the monitor.

Dim Image If this is a problem with the monitor itself, it's most likely caused by a failing backlight. However, it could also be a configuration setting. On an external monitor, open the onscreen menu and try adjusting the brightness. On a laptop, use the function keys to adjust the brightness. Figure 13.6 shows a laptop keyboard where the keys are F9 and F10. (Back in Figure 13.1, they are F5 and F6 on the top example, and F11 and F12 on the bottom one.) If no amount of adjustment makes the image bright enough for normal use, then replace the backlight.

FIGURE 13.6 Laptop function keys

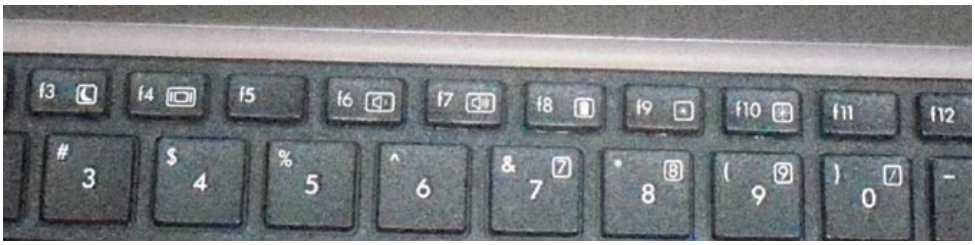

Other Display Issues

In this final section on video, projector, and display issues, we look at a few symptoms that don't fit nicely into our other sections. This is kind of the grab bag of random video and display problems.

Audio issues and intermittent projector shutdown are the final two symptoms you need to know for Exam Objective 5.3.

Audio Issues Many display units today have built-in speakers. The most common reason people have audio problems is because something is muted, but it could also be a cable or connection issue. HDMI cables carry audio, but DVI cables do not.

First, check the display unit to ensure it's not muted and that the volume is turned up to a reasonable level. This is done on the display's onscreen menu. Next, check to ensure that the computer's audio output is set to the correct device and that it's not muted or the volume isn't turned down. In Windows, right-click the speaker icon on the taskbar, then choose Open Sound Settings (Figure 13.7). (If there is no speaker icon to click, go

to Start ➢ Settings ➢ System ➢ Sound.) Choose the correct device in the Output section, and ensure that the master volume is turned up.

FIGURE 13.7 Windows sound settings

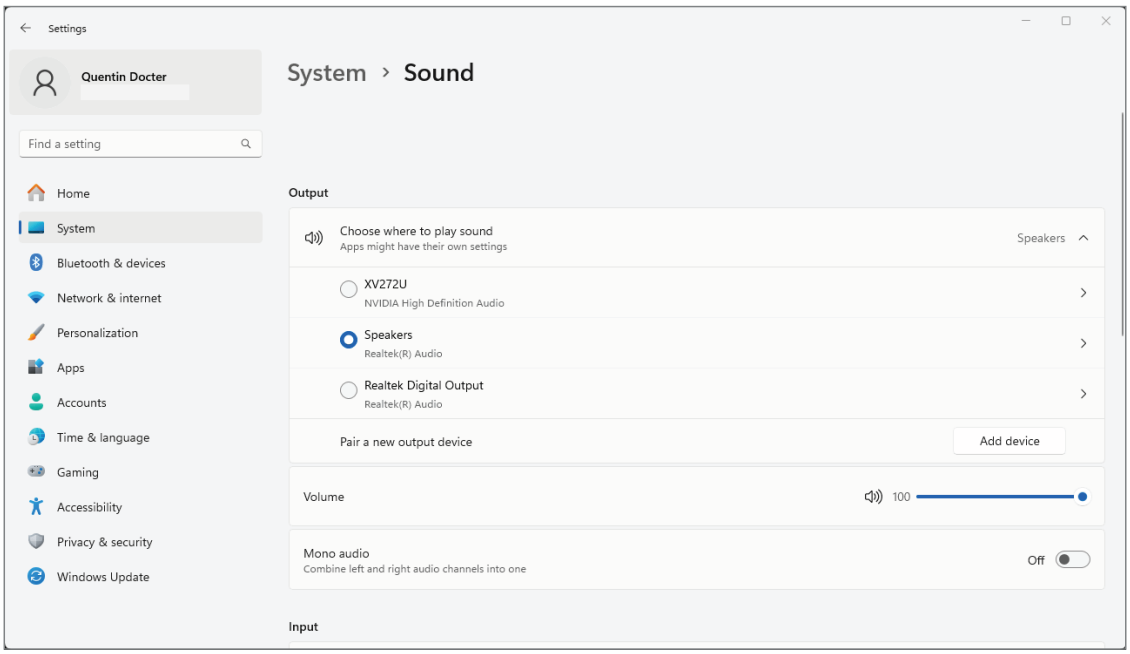

Intermittent Projector Shutdown We noted earlier that projectors create a lot of heat. When a projector overheats, it will shut itself off to avoid frying components or the bulb. This is the most likely cause of intermittent shutdowns. After the projector cools off, perform a little maintenance cleaning. Most projectors have an air filter to keep dust and debris out of it—check to ensure that's clean and replace it if necessary. Also check to make sure the cooling fan is operational and blowing out warm air.

Monitors can shut down intermittently as well due to overheating. It was more common on older CRT monitors than it is on LCD ones, but it can still happen. Be sure the air vents on the back of the display unit are clear from dust and debris. If the problem persists, it's best to replace the monitor.

Booting into Low-Resolution Mode Sometimes an operating system refuses to boot into its normal mode and instead boots into a basic video mode. Newer versions of Windows call this "low-resolution" mode, whereas older versions called it VGA mode. This is normally caused by one of two things. One, the video card was set to a resolution it can't handle, or two, there's something messed up with the video card driver. After the OS boots, try reducing the resolution and reboot. If it's still a problem, reinstall the video card driver. If new drivers don't help, the video card may be failing.

Other graphics issues can be attributed to the memory installed on the video card. This is the storage location of the screens of information in a queue to be displayed by the monitor. Problems with the memory modules on the video card have a direct correlation to how well it works. It follows, then, that certain unacceptable video-quality issues, such as jerky refresh speeds or lags, can be remedied by adding memory to a video card (if possible). Doing so generally results in an increase in both quality and performance. If you can't add memory to the video card, you can upgrade to a new one with more memory.

Troubleshooting Common Mobile Device Issues

Mobile devices, for the most part, are essentially the same types of devices as desktops, but troubleshooting the two can feel very different. While the general troubleshooting philosophies never change—steps such as gathering information, isolating the problem, and then testing one fix at a time—the space and configuration limitations can make troubleshooting smaller devices more frustrating. For purposes of the discussion here, the term *mobile devices* includes laptops and anything smaller.

We will look at five areas where mobile devices could have different problems from their desktop counterparts: power and heat, input/output, connectivity, damage, and other symptoms. Much of what we cover will be more closely related to laptops than smaller mobile devices, but the concepts generally apply to mobile computers of all sizes. We'll call out specifics for small mobile devices where applicable.

Working on Laptops and Small Mobile Devices

For whatever reason, it's easier to lose screws and other small pieces when working on laptops than it is with desktops. Don't forget these key concepts when working on laptops:

- Refer to the manufacturer's documentation.
- Use the appropriate hand tools.
- Have a clear organization method for your parts and screws.
- Document and label screw and cable locations. Take photos if you can. For a review of these four concepts, see Volume 1 Chapter 4, "Laptop and Mobile Device Hardware."

A few of the mobile device issues in Exam Objective 5.4, "Given a scenario, troubleshoot common mobile device issues," are very similar to the A+ 220-1202 Exam Objective 3.2, "Given a scenario, troubleshoot common mobile OS and application issues," and 220-1202 Exam

> Objective 3.3, "Given a scenario, troubleshoot common mobile OS and application security issues." The implication is that for the A+ 220-1201 exam you will be asked about hardware issues and for the A+ 220-1202 exam it would be OS related. They're similar and often intertwined, so you may see some similar content in Volume 2 Chapter 7. We will do our best to separate out the hardware angle from software-related issues, but as we have stated before, sometimes it's a fine line as to what's really causing the problem.

Power and Heat Issues

Mobile devices are different from desktops in that they're designed to work without a continually plugged-in power source. That freedom introduces complexities that can cause power-related problems, though—specifically, the battery and charging the battery. In addition, because of their compact nature, mobile devices are more prone to overheating. In this section we'll look at battery-related issues as well as overheating problems.

Battery and Charging Issues

Mobile devices are of course meant to be mobile and not plugged in at all times. It's a bit ironic, then, that a good question to ask if a mobile device doesn't seem to power up is, "Is it plugged in?" Everyone hates getting asked that question, but it's a critical question to ask, even with mobile devices. If the device works when it's plugged in but not unplugged, you've narrowed down the problem. You can't assume that the battery is working (or is attached) as it's supposed to be. Always check power and connections first!

Poor battery health, improper charging, and swollen battery are symptoms to know as part of Exam Objective 5.4.

Poor Battery Health

If the laptop works while it's plugged in but not while on battery power, the battery itself may be the culprit. A battery problem might also prevent the laptop from powering on when plugged into an electrical outlet. In either of these cases, the first fix to try is to disconnect the battery from the motherboard, let it sit for a minute or two, and then plug it back in.

As batteries get older, they are not able to hold as much of a charge and, in some cases, are not able to hold a charge at all. That is to say, the battery health may be poor. If the battery won't charge while the laptop is plugged in, try removing the battery and reinserting it. If it still won't charge, you might want to replace the battery.

Another issue that small devices can have is an extremely short battery life. We're not talking about when people complain that their laptop only runs for an hour and a half when they are playing a Blu-ray while surfing the Internet and talking to their friends on their Bluetooth headset. No, that's bound to drain your battery quickly. What we're referring to here is when a laptop battery only lasts for an hour or so after a full charge with normal

usage, or if a mobile phone battery is only able to power the device for 30 minutes or so. These things happen.

If it's a laptop, you can try to perform a battery calibration, as we discussed in Volume 1 Chapter 4. For all mobile devices, you can try to drain the battery completely and then charge it fully before turning the device back on. If these options don't work, then it's likely that the battery needs to be replaced.

Improper Charging

Many laptop power adapters have a light indicating that they're plugged in. If there's no light, check to make sure that the outlet is working, or switch outlets. Also, most laptops have a power-ready indicator light when plugged into a wall outlet as well. Check to see if it's lit. If the outlet is fine, try another power adapter. They do fail on occasion.

Smaller mobile devices will have a lightning bolt next to their battery icon or an animated filling battery when charging. If the device doesn't appear to charge, the same culprits apply: it could be the outlet, the adapter, or the device itself.

If you're working on a DC adapter, the same concepts apply. Check for lights, try another adapter if you have one, or try changing plugs, if possible. For example, if you're using a DC outlet in a car, many newer models have secondary power sources, such as ones in the console between the seats.

Another thing to remember when troubleshooting power problems is to remove all external peripherals. Strip your laptop down to the base computer so that there isn't a short or other power drain coming from an external device.

Windows has built-in power management features to help conserve laptop battery life. In Windows, open the Power applet in Settings (go to Start ➢ Settings ➢ System ➢ Power.) Once there, you can configure different power-saving settings to maximize the battery life of your laptop.

Swollen Battery

The last power issue that we need to discuss is a *swollen battery*. As the term suggests, the battery physically swells in size. It can be caused by a number of things, including manufacturer defects, age, misuse, using the wrong adapter for charging, or leaving the laptop constantly plugged into a wall outlet. Inside the battery, the individual cells become overcharged, causing them to swell. Sometimes the swelling is barely noticeable, but it can cause the device case to crack or pop apart. Other times it's pretty obvious, such as the one shown in Figure 13.8.

If you have a swollen battery, turn the device off immediately and make sure that it's not plugged into a charger. If the battery is removable, you can try to remove it, if you wish, but be very careful. Swollen batteries are more prone to explosions than normal batteries because the casing is already compromised. If you are able to remove it, place it into a safe container, just in case there are further issues. If the battery is not removable, it's time for a new device. In either case, take the battery or device to a proper recycling center to dispose of it. Never just throw it in the trash; it can explode and harm sanitation workers, as well as cause significant damage to the environment.

FIGURE 13.8 iPhone with a swollen battery

Source: Mpt-matthew / Wikimedia Commons / CC BY SA 3.0

Overheating

Smaller devices have a greater potential to overheat than do their larger cousins. Space is at a premium, so all the components are packed tightly together, which means less room for each component to breathe. Manufacturers realize this, of course, so they use components that generate less heat. Overheating can still be a problem, though. If your mobile device is overheating, turn it off to let it cool down. It could be from overuse, or perhaps it did not have proper ventilation (for example, if it was stuffed into someone's pocket or purse). On laptops, check to ensure that the cooling fan is working and not full of dust or debris. If the overheating is persistent, you have a few options. The first is to test or replace the battery, as that's the most likely culprit. If overheating still happens, you may have to replace the device.

Exam Objective 5.4 includes overheating as a symptom to know.

Input Problems

Laptop keyboards aren't as easy to switch out as desktop keyboards. You can, however, very easily attach an external USB keyboard to your laptop if the keys on your laptop don't appear to work.

If the keyboard doesn't seem to respond at all, try pressing the Num Lock and Caps Lock keys to see if they toggle the Num Lock and Caps Lock lights on and off. If the lights don't respond, the keyboard isn't functioning properly. Try rebooting the system. (You will probably have to press and hold the power button for 5 seconds, and the system will shut off. Wait 10 seconds, and press the power button again to turn it back on.) If that doesn't fix the problem, you probably have faulty hardware.

Another problem unique to laptop keyboards is the *Fn key*. (It can be your friend or your enemy.) You can identify it on your laptop keyboard because it's in the lower-left corner and has the letters *Fn* on it (often in blue), as shown in Figure 13.9. If the Fn key is stuck on, the only keys that will work are those with functions on them. If you look at other keys on your laptop, several of them will have blue lettering too. Those are the functions that the keys may perform if you press and hold the Fn key before pressing the function key that you want. If the Fn key is stuck on, try toggling it just as you would a Caps Lock or Num Lock key.

FIGURE 13.9 The Fn key on a laptop

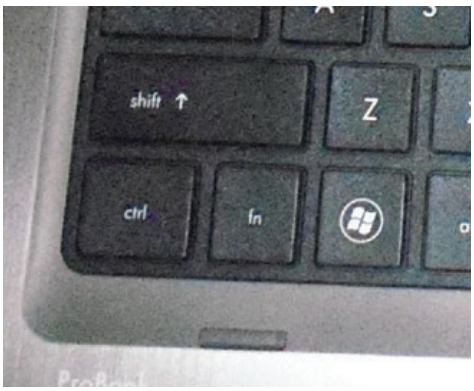

If another key on your laptop keyboard is stuck, you need to determine if the contact is having problems or if the key itself is stuck. If the key is not physically stuck but the laptop thinks it is, rebooting generally solves the problem. If the key physically sticks, you can try blowing out underneath the key with compressed air, or use a cotton swab slightly dampened with water (or isopropyl rubbing alcohol) to clean underneath the key. Make sure to clean the entire surface underneath the sticking key. If none of this resolves the issue, you might need to replace the keyboard.

One of the conveniences that users often take advantage of in laptops is a built-in pointing device. Most laptops have touchpads or point sticks that function much like a mouse. They're nice because you don't need to carry an external mouse around with you. While these types of devices are usually considered very handy, some people find them annoying. For example, when you are typing your palm might rest on the touchpad, causing erratic pointer behavior. This is referred to as a *ghost cursor* because it seems like the cursor just randomly jumps all over the screen. You can turn the touchpad off or calibrate touch settings in Windows 11 through the Touchpad app or Mouse settings app (Volume 1 Chapter 4, Exercise 4.8 shows how to do that). While understanding that you can turn it off on purpose, remember that it can be turned off accidentally as well. Check to make sure that it's enabled. Some laptops allow you to disable or change the sensitivity of the touchpad as well, just as you can adjust the sensitivity of your mouse.

Input symptoms in Exam Objective 5.4 that you should know include digitizer issues, cursor drift/touch calibration, and stylus does not work.

Another potential issue is *cursor drift*, where the mouse cursor will slowly drift in one direction even though you are not trying to make it move. This issue is generally related to the point stick not centering properly after it's been used. If you have cursor drift, try using the point stick and moving it back and forth a few times to get it to re-center itself. You can also try calibrating it within the operating system (most manufacturers make it a tab in Mouse properties), or rebooting. If the problem persists, either disable or replace the point stick.

Finally, here are three issues you may encounter with mobile device displays: digitizer issues, stylus does not work, and broken screens. Recall from Volume 1 Chapter 4 that a digitizer is a device that can be written or drawn on, typically with the touch from a human finger or a stylus. Most mobile devices have a digitizer built into the display unit. It may be the glass of the display itself, or it might be implemented as an overlay for the display. Either way, if it's not functioning, either by not registering a finger's touch or that of the stylus, that can cause problems. You can probably work around it on a laptop, but a smartphone or tablet with a nonworking digitizer is pretty useless.

In touch-enabled Windows devices, you can calibrate digitizers by going to Control Panel ➢ Tablet PC Settings ➢ Calibrate (under the Display tab). Rebooting may also help. For iOS and Android tablets and phones, if the digitizer isn't working the only troubleshooting step is to power it off and restart the device. With any device, if the digitizer isn't working, the next step is to replace the screen or the device.

If a powered stylus isn't working, it could be that the battery in it is dead. Check that first. You may also be able to replace the stylus tip, depending on the model. Finally, you can restart the mobile device, remove a screen protector (if applicable), and check to ensure the stylus is paired properly.

A broken screen, while unfortunate, is all too common with mobile devices. Considering the beating they take on a regular basis, it's a little surprising it doesn't happen more often. First, to help avoid broken screens, make sure all of your mobile users have screen protectors and drop-resistant cases. If a screen does get broken, either replace the screen or replace the device.

Connectivity Issues

Nearly every mobile device sold is equipped with integrated wireless networking, and most have Bluetooth built in as well. In many cases, the wireless antenna is run into the display panel. This allows the antenna to stand up higher and pick up a better signal.

Poor/no connectivity is a common symptom listed in Exam Objective 5.4.

If your wireless networking isn't working on a laptop, do the following:

1. Ensure that the network card is turned on. Many laptops have an external switch or button on the front or side or above the keyboard that can toggle the network card on and off. In some cases, it will be a function key. Be sure that this is set to the On position. Figure 13.10 shows a toggle above the keyboard. (It's the one on the left that looks like an antenna.)

FIGURE 13.10 Network card toggle switch above the keyboard

2. Next, make sure the wireless card is enabled through Windows. You generally do this in Windows 11 as follows:
   a. Open Start ➢ Settings ➢ Network & Internet.
   b. Click Advanced Network Settings.
   c. Check to ensure that the Wi-Fi connection is enabled. If not, click the Enable button.

   Some network cards have their own proprietary configuration software as well. You can also often check this by clicking a tab (often called Wireless Networks, in the network card's Properties page) to see if you're getting a signal and, if so, the strength of that signal.
3. Check the strength of the signal.

   A weak signal is the most common cause of intermittent wireless networking connection problems. If you have intermittent connectivity and keep getting dropped, see if you can get closer to the wireless access point (WAP) or remove obstructions between you and the WAP. Failing network cards and connectivity devices can also cause intermittent wireless networking connection failures.

If you have a USB network adapter, try unplugging it and plugging it back in. Make sure that Windows recognizes the card properly.

If the wireless connection fails but the system has a wired RJ45 port, try plugging it in. For this, you will need an Ethernet cable and, of course, a wired network to plug it into. But if you get lights on the NIC, you might get on the network.

Real World Scenario

Potential Wireless and Wired Conflicts

A short time ago, a friend of ours was frustrated because he couldn't get to the network in his office with his laptop plugged into his docking station. He had used the laptop at home the night before and gotten on his wireless network without a problem. But this day, his wired connection would not work. He checked his cables (always your first step) and saw that there were lights (a good sign). He had tried to access both the Internet and intranet sites, but to no avail.

We opened a command prompt and ran `ipconfig`. He didn't have an IP address, but we noticed that his built-in wireless card was listed and active.

What he needed to do was to disable his built-in wireless card. He had enabled the wireless to work at home, and it was still enabled. Because it was enabled, the wireless card was trying to obtain an IP address, and it refused to let the wired "portion" of the card pick up an address from the company DHCP server. (There was no wireless in the building due to strict security policies.) After he disabled his wireless card, his wired connection picked up an IP address, and all was well.

Many laptop network cards have a wired connection in addition to their wireless capabilities. For some of them, the wired connection will not work if the wireless is enabled. It's an attempt to prevent conflicts if both connection types are active.

The principles behind troubleshooting network or Bluetooth connectivity issues on mobile phones and tablets are the same as on laptops. The big difference is that you can't try an external network card if your internal card is failing. We originally looked at some of these settings in Volume 1 Chapter 10, "Mobile Connectivity and Application Support," but now is a good time to review them. The first thing to check is that the network connection or Bluetooth is enabled, which also means double-checking that airplane mode is not turned on. On Android and iOS devices, this is done through Settings. Figure 13.11 shows iOS network settings and Android network connections. Toggle the connection off and then back on to reset it; often, that will resolve connectivity issues.

FIGURE 13.11 iOS (a) and Android (b) network settings

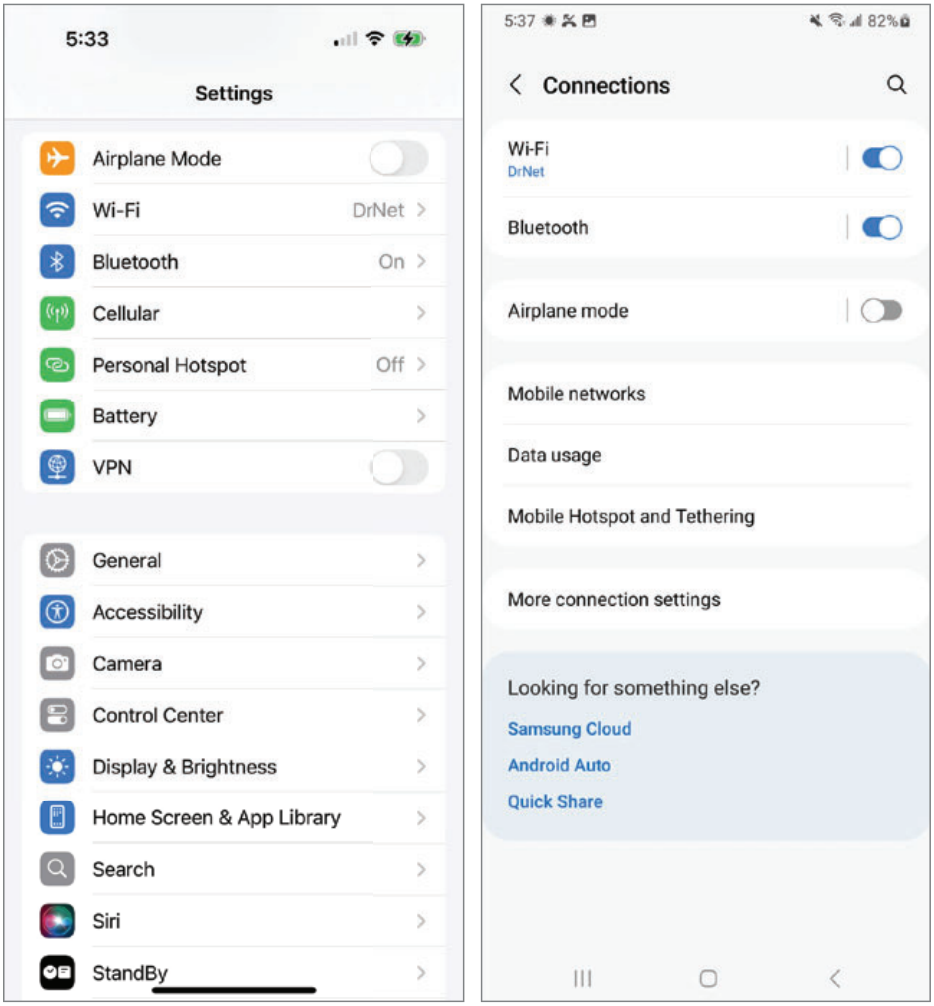

Another way to access network settings in iOS is from the Control Center. You can do this from both the lock screen and the home screen. Simply swipe your finger down from the very top of the iPhone's touchscreen, and you will get the Control Center, similar to what's shown in Figure 13.12. In Android, open the notifications area by swiping down from the top of the screen (you may need to swipe down twice), and network settings will be there as well, also shown in Figure 13.12.

Poor or intermittent connectivity on mobile devices is usually a function of the distance from the access point. Wi-Fi has maximum ranges that can be dramatically shortened by obstructions such as steel or cinder-block walls or interference from fluorescent lights, motors, or microwaves. Usually, moving closer to the access point will resolve these issues.

FIGURE 13.12 iOS Control Center (a) and Android notifications area (b)

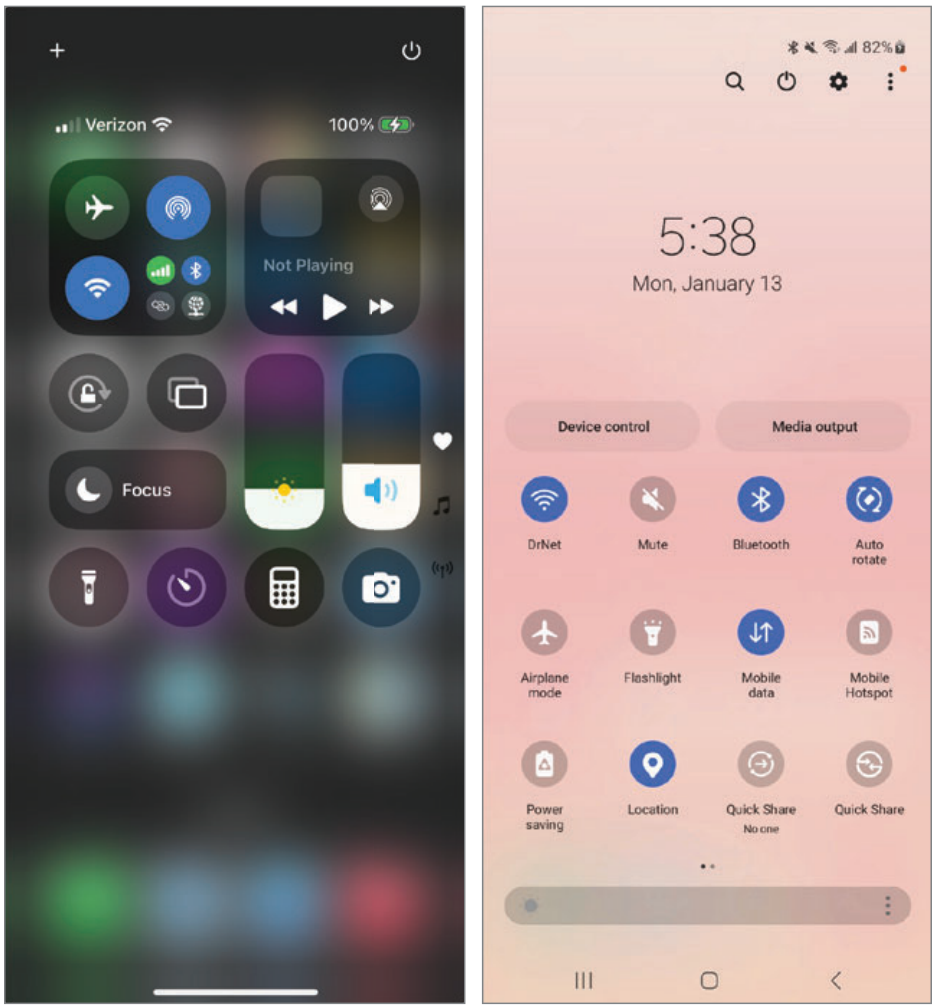

Physical Damage and Malware

Mobile devices take much more of a beating than stationary devices do, which is why cases and screen protectors are needed accessories. Sometimes things happen, though, and a device gets physically damaged. Here we will look at two types of physical damage: liquid damage and physically damaged ports. Then we will take a look at malware and how to avoid issues it can cause.

Exam Objective 5.4 lists liquid damage, physically damaged ports, and malware as symptoms you should be able to troubleshoot.

Physical Device Damage

A device can become damaged in any number of ways, with dropping being the most common. Even if you have a great case on your phone, an airborne expedition down a flight of concrete stairs probably isn't going to have a happy ending. Similarly, liquid can do nasty things to electronics as well.

If a laptop gets doused in a liquid, it's best to turn it off as soon as possible and let it dry out. If it was a spill on the keyboard and no liquid got into ports, the computer is probably salvageable. If liquid got inside the ports, they could be very difficult to clean out and the liquid can cause connectivity problems for those ports.

For significant liquid damage, the laptop can be taken apart and the components—even circuit boards such as the motherboard—can be cleaned with demineralized water and a lint-free cloth. Disassemble the components, clean them with water and the cloth, and let them thoroughly dry. Reassemble, and see if it works. Several years ago, one of our friends had a toddler son who decided to relieve himself on the laptop. The demineralized water and a thorough drying did the trick—the machine still worked. We can't guarantee success in every situation, however!

Mobile devices are much more liquid-friendly than laptops are. Many phones and tablets today are water-resistant, if not entirely waterproof. Every smartphone that is considered water-resistant will have an ingress protection (IP) rating, such as IP67 or IP68. The first digit, which will be from 0–6, represents the device's ability to withstand solid foreign material such as dust. The second digit, which will be from 0–8, shows its moisture resistance. Sometimes you will see a rating such as IPX6, which means the device has not been tested for dust resistance and has a moisture resistance rating of 6.

IP ratings were created by the International Electrotechnical Commission (IEC) to provide standard guidance as to the degrees of protection provided by enclosures of electrical equipment.

For a device to be considered waterproof, it needs to have a moisture resistance rating of 7 or 8. A 7 rating means the device is protected from damage from immersion in water with a depth of up to 1 meter (3.3 feet) for up to 30 minutes. An 8 rating is given to devices that can withstand greater depth and time of immersion, which must be specified by the manufacturer. For the best in protection, buy a device with an IP67 or IP68 rating. Devices with no IP code might or might not survive heavy rain, sprays from water sources, or an accidental dunking.

If you suspect a mobile device has suffered water damage, first, immediately turn it off. Remove the case and anything else that can be removed, such as the SIM card and possibly the battery. Dry everything you can with a lint-free cloth. Then you have a few choices. One is to let it air dry for at least 48 hours. Or you can try the "rice trick." That is, put the device in a sealed container, covered in uncooked rice, and let it sit for 48 hours. The rice should soak up all of the water in the device. Note that professional opinions are divided on the rice trick. Some experts say it works well, others say that residue in the rice can damage electronic components. But if the phone isn't working anyway, how much risk is it to try?

When it comes to ports and damage, at least one of two things is always true: either the ports are obviously damaged, or they simply fail to work. In either case, the only remedy is to replace the port, which usually means replacing several components, including the motherboard on a laptop, or replacing the entire mobile device.

Malware

Malware is malicious software designed to damage, disrupt, or gain unauthorized access to a computer system. Malware infections are one of the most common security risks that you will encounter. Let's look at malware on laptops as well as mobile devices.

Laptops and Malware

Laptops running Windows or macOS have the same vulnerability to malware as their desktop cousins do. To help guard against malware, install antimalware software. These are the four main classes of applications to help protect your system against malware and hackers:

- Antivirus software defends against viruses, worms, and Trojan horses.
- Antispyware software defends against adware and spyware.
- Antispam software reduces the amount of junk email you receive.
- Software firewalls block potentially dangerous network traffic.

There are also suites available that combine multiple security functions; for example, the Norton security products can include antivirus, antimalware, and antispam features, along with identity protection tools, a software firewall, a backup tool, and a PC tune-up tool. The level of coverage depends on the product you purchase. McAfee offers similar solutions. In addition, there is some overlap between the types of threats each application guards against; for example, an antivirus program may also target some types of nonvirus malware.

Even if you have an antimalware application installed, it's not perfect. Occasionally a virus or other malware may get around it, such as a new to the world threat, or even an older one if you haven't updated your definitions lately. When a system is infected with a virus, a worm, a Trojan horse, or other malicious software, you need to remove it immediately. Here are the five steps to take to remove malware:

1. **Identify malware symptoms.** Before doing anything major, it is imperative first to be sure you are dealing with the right issue. If you suspect malware, then try to identify the type (spyware, virus, and so on) and look for the proof needed to substantiate that it is indeed the culprit.
2. **Quarantine the infected system.** Once you have confirmed that malware is at hand, then quarantine the infected system to prevent it from spreading the malware to other systems. Bear in mind that malware can spread any number of ways, including through a network connection, email, and so on. The quarantine needs to be complete to prevent any spread.
3. **Remediate infected systems.** The steps taken here depend on the type of malware with which you're dealing, but it should include updating antivirus software with the latest definitions and using the appropriate scan and removal techniques.

4. **Schedule scans and updates.** The odds of the system never being confronted by malware again are slim. To reduce the chances of it being infected again, schedule scans and updates to run regularly.
5. **Educate the end user.** Education should always be viewed as the final step. The end user needs to understand what led to the malware infestation and what to avoid, or look for, in the future to keep it from happening again.

The five malware removal steps are listed in A+ Exam 220-1202 Objective 2.6, and they are also covered in Volume 2 Chapter 7. We put them here just in case they appear in Exam 220-1201 questions as well.

Mobile Devices and Malware

Many people believe, incorrectly, that mobile devices running iOS or Android can't get infected with malware. Others believe that iOS is totally safe but Android is less so. The second statement is a little closer to the truth, but iOS is still vulnerable to malware. Let's take a look at the four most common ways mobile devices can contract malware:

Malicious Apps The most common way that malware ends up on mobile devices is through downloading malicious apps. Downloading apps from official sources such as Google Play or the App Store is almost always safe. Apps are vetted and should be good, although Apple is a bit more stringent on developer guidelines than Google is. If someone has jailbroken their phone and is downloading apps from random places though, this is a great way to contract a malware infection.

Jailbreaking is when someone removes Apple's security restrictions, allowing them full access to the device. This means the user can install unauthorized apps and make changes Apple would not otherwise approve of. Jailbreaking is illegal in some countries, such as the United States, for violating copyright laws. The equivalent to jailbreaking on Android is called *rooting*. Both of these procedures can weaken device security and usually voids the device's warranty.

Operating System Vulnerabilities If you have a mobile device, it's guaranteed that you've seen the notification to update your OS to a newer version. One of the reasons for that is because developers may have uncovered a vulnerability that hackers can exploit. A general rule of thumb is to always update to the newest version of the mobile operating system that you can.

Suspicious Links This is the same as it would be on a laptop or a desktop. If you get an email with a suspicious link, don't click it. Clicking it can install malware on your device without your knowledge.

Unsecure Networks This, too, is the same as it would be for laptops. Using an unsecure network doesn't really pose much of a risk for you downloading malware, but it could easily allow a hacker to read the data transmitted to or from your device. A VPN can help eliminate these issues on unsecured public networks.

A few good rules of practice to avoid malicious apps or malware on a mobile device are:

- Don't jailbreak mobile devices. This compromises built-in security mechanisms.
- Download apps only from official app stores.
- Always update the OS to the most current version.
- Avoid clicking links if you are unsure of the source.
- Use a VPN when on public networks.

Antimalware software such as Norton and Avast can be purchased for iOS and Android as well.

Other Mobile Device Issues

The last two issues we need to cover don't quite fit neatly into any of the categories we've presented so far. They are:

- Unable to install new applications
- Degraded performance

Exam Objective 5.4 common symptoms include unable to install new applications and degraded performance.

Unable to Install New Applications

Apps are critical components to any personal computing system. Without them, the block of hardware in your hand or lap isn't very useful. So it stands to reason that not being able to install new ones is problematic.

Let's tackle laptops first. From a software standpoint, laptops and desktops are basically the same. That is, they both rely upon hard drive space for long-term storage, and the OS coordinates the app interaction with hardware and controls security. Therefore, all of these could be considered possible culprits if a new app won't install. Sometimes you will get an error message that leads you to the solution, whereas other times it just won't install. Or, it will say that it's installed but the app refuses to open.

If the app simply won't install, here are some things to check:

- Is there enough hard drive space?
- Is the app compatible with the operating system?

- Does the user account trying to install the app have sufficient privileges? If not, try installing while logged in as an administrator.
- Are there security settings prohibiting the installation? One common blocker is security software such as an antivirus suite.
- Is the installation file corrupted? It might be worth downloading it again.

Of course, as mentioned previously, if you get an error message that can go a long way towards solving the problem. If not, and hard drive space is not an issue, then the most common cause is security software.

If the app installs but refuses to run, the list of potential causes is shorter. It's most likely one of these:

- The installation file was corrupted.
- The user account doesn't have the permissions needed to run the app.
- Insufficient memory.

The last two are usually accompanied by an error message. If the install is corrupted, it's a toss-up on if there will be a message or not. Reboot the system and try the app again. If it still doesn't work, delete and reinstall.

Smaller mobile devices such as tablets and smartphones have a similar list of potential causes, with a few unique twists. Here's what to check:

- Is there enough storage space?
- Is the OS the most current version?
- Does the user account have permission to install the app?
- Is the Internet connection strong and stable?
- Is the app store the most current version?
- Within the app store, is the user account and payment information correct?
- Is the app an official one from the app store? iOS in theory won't let you install an app that's not from the app store. On Android devices, you can, but you need to enable installation from "unknown sources" first.

If an app appears to install but doesn't work, hopefully you know the answer by now. Restart the device first and then ask questions. If after restarting the app still doesn't work, and the user has appropriate permissions, the solution is to remove the app and reinstall.

Degraded Performance

As devices get older, their performance will naturally start to decline. Electronic components subjected to heating and cooling, over time, will simply wear out. In addition, the term *degraded performance* can be very open to interpretation. Some users will become very unhappy when they perceive that their device is two milliseconds slower than it used to be.

What we mean by degraded performance is when the device is noticeably and demonstrably showing signs of slowing down. This can include shorter battery life, slow apps, sluggish responsiveness to input, consistently poor wireless connections, and other symptoms. Of course, try restarting the device to see if that helps. If the device has overheated, leave it off for a while before turning it back on.

Sometimes performance gradually slows over time, and other times it sharply declines. In the latter case, it's usually due to some physical damage or a new app that took up too much storage space. The short list of common causes of degraded performance is:

- Physical damage.
- Age. Older devices will slow down.
- Lack of memory or storage space. Devices with less than 1 GB of free space might experience slowdowns.
- Failing or old battery.
- Outdated OS or apps.

By now, you should have a good idea of how to solve the listed problems. Restart, uninstall apps, update the OS and apps, delete files (or photos, videos, and music), and replace hardware if needed and possible. If all else fails, replace the device.

Real World Scenario

Batterygate

For many years, it was rumored that Apple intentionally slowed down older devices. The theory was that Apple did so to increase sales of newer devices because the performance on old ones became untenable for many users.

In 2017, Apple admitted to intentionally throttling down the speed on devices with older batteries in what became known as the "Batterygate" controversy. Their reason was that it would help conserve battery life, as well as prevent unexpected shutdowns on those devices. Apple was subsequently fined hundreds of millions of dollars by several agencies for failing to disclose this practice to its customers.

Since then, Apple has introduced more transparent battery health information in iOS. This can be found in Settings ➢ Battery, shown in Figure 13.13. Tapping on Battery Health & Charging (also Figure 13.13) shows battery options. If the device is being throttled, there will be a message in the Peak Performance Capability section, with a Performance Management option for the user to disable throttling.

FIGURE 13.13 iOS Battery (a) and Battery Health & Charging (b) settings]

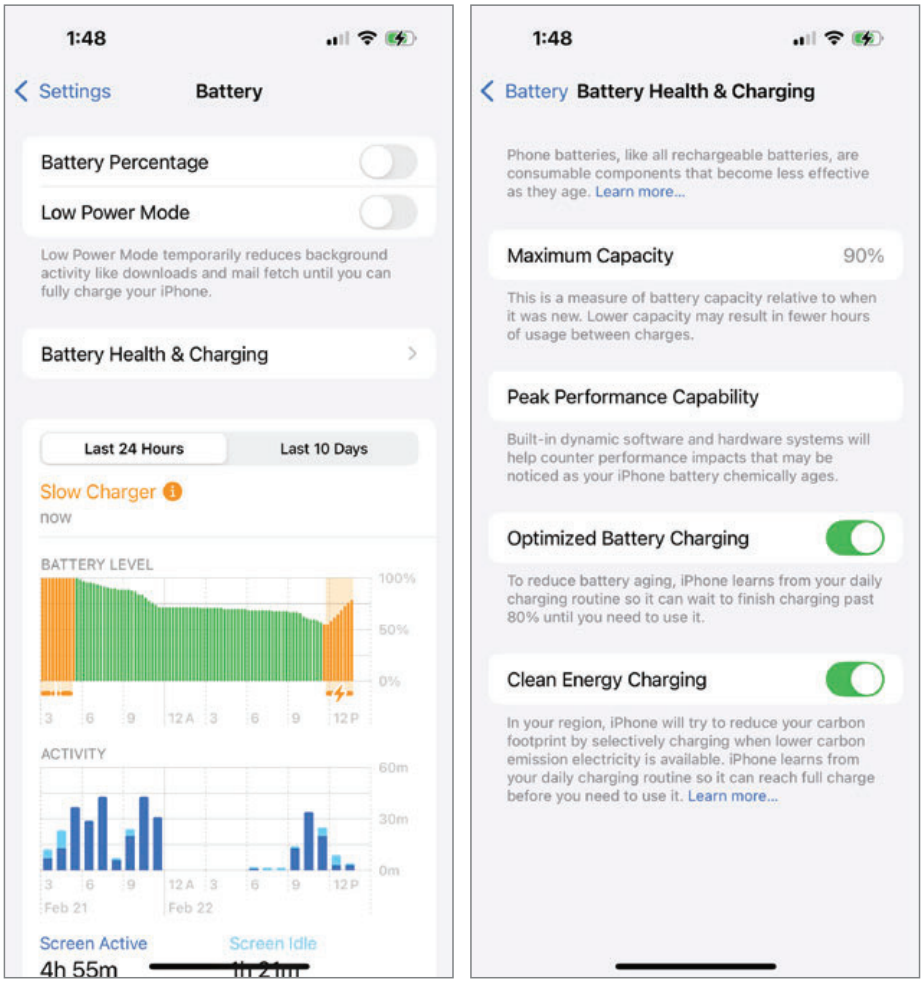
1:48
Settings
Battery
Battery Percentage
Low Power Mode
Low Power Mode temporarily reduces background activity like downloads and mail fetch until you can fully charge your iPhone.
Battery Health & Charging
Last 24 Hours
Last 10 Days
Slow Charger
now
BATTERY LEVEL
100%
50%
0%
3 6 9 12 A 3 6 9 12 P
ACTIVITY
60m
30m
0m
3 6 9 12 A 3 6 9 12 P
Feb 21
Feb 22
Screen Active
4h 55m
Screen Idle
1h 21m
1:48
Battery
Battery Health & Charging
Phone batteries, like all rechargeable batteries, are consumable components that become less effective as they age. Learn more...
Maximum Capacity
90%
This is a measure of battery capacity relative to when it was new. Lower capacity may result in fewer hours of usage between charges.
Peak Performance Capability
Built-in dynamic software and hardware systems will help counter performance impacts that may be noticed as your iPhone battery chemically ages.
Optimized Battery Charging
To reduce battery aging, iPhone learns from your daily charging routine so it can wait to finish charging past 80% until you need to use it.
Clean Energy Charging
In your region, iPhone will try to reduce your carbon footprint by selectively charging when lower carbon emission electricity is available. iPhone learns from your daily charging routine so it can reach full charge before you need to use it. Learn more...

Troubleshooting Printer Issues

Even though society is moving away from paper forms, printers are still very common peripherals. Printers are also the most complex peripheral as far as troubleshooting is concerned; this arises from complications in putting ink to paper. There are several different ways that this can be accomplished, but the end result is all pretty much the same.

Different types of printers work in different ways, so you would expect that laser printers might have different issues from impact printers. Because problems are often dependent on the type of printer you're using, we've chosen to break down this discussion by printer type. We'll start with a quick review of the technology and then get into specific issues. At the end,

we'll look at the process of managing the print spooler, which is the same regardless of the printer type in use.

Printer manufacturer websites are great places to look to find troubleshooting information. They often provide descriptions of problems and detailed instructions for resolving the issue. Most printers also come with management software that you can install on your computer, which may be able to assist you in troubleshooting any issues that you have.

Help! I Can't Install My Printer!

Before you can print, you must, of course, install the printer on the computer you want to print from. Sometimes, though, installing the printer can be problematic. The first common issue you will run into is that the printer isn't found during the installation process. If this is the case, check to make sure that the printer is on and connected to either the computer or the network. Wireless network printers are notorious for not being "found." If needed, re-seat the connections and try again, or see if you can connect the printer using a different method (for example, USB instead of wireless).

The second common issue is that the printer driver won't install. Be sure that any other instances of that type of printer aren't already installed on the computer. If you just removed the old printer driver, reboot before reinstalling the new one. And, finally, download and install the newest driver from the manufacturer's website.

Impact Printer Issues

Impact printers are so named because they rely on making a physical impact in order to print. These are typically dot-matrix or daisy wheel printers. The impact printer's printhead will be activated and forced up against the ink ribbon, making a character or pattern on the paper. Impact printers are relatively simple devices; therefore, only a few problems usually arise. We will cover the most common problems and their solutions here.

Low Print Quality

Problems with print quality are easy to identify. When the printed page comes out of the printer, the characters may be too light or have dots missing from them. Table 13.1 details some of the most common impact print quality problems, their causes, and their solutions.

TABLE 13.1 Common impact print quality problems

Characteristics	Cause	Solution
Consistently faded or light characters	Worn-out printer ribbon	See if you can adjust the printhead to be closer to the ribbon. If not (or if it doesn't help), replace the ribbon with a new, vendor-recommended ribbon.
Print lines that fade from dark to light as the printhead moves across the page	Printer ribbon-advance gear slipping	Replace the ribbon-advance gear or mechanism.
A small, blank line running through a line of print (consistently)	Printhead pin stuck inside the printhead	Replace the printhead.
A small, blank line running through a line of print (intermittently)	A broken, loose, or shorting printhead cable	Secure or replace the printhead cable.
A small, dark line running through a line of print	Printhead pin stuck in the out position	Replace the printhead. (Pushing the pin in may damage the printhead.)
Printer making a printing noise, but no print appears on the page	Worn, missing, or improperly installed ribbon cartridge	Replace the ribbon cartridge correctly.
Printer printing garbage, such as garbled characters	Cable partially unhooked, wrong driver selected, or bad printer control board (PCB)	Hook up the cable correctly, select the correct driver, or replace the PCB (respectively).

Exam Objective 5.6, "Given a scenario, troubleshoot printer issues," lists several symptoms to know for the exam. Table 13.1 covers three of them for impact printers:

- Lines down the printed pages
- Garbled print
- Faded prints

These symptoms will appear in other printer type sections too; of course, the solution for each will depend on the type of printer causing the problem.

Printout Jams Inside the Printer

Printer jams (aka "the printer crinkled my paper") are very frustrating because they always seem to happen more than halfway through your 50-page print job, requiring you to take time to remove the jam before the rest of your pages can print. A paper jam happens when something prevents the paper from advancing through the printer evenly. There are generally three causes of printer jams: an obstructed paper path, stripped drive gears, and using the wrong paper.

Paper jams, paper not feeding, and multipage misfeed are issues listed in Exam Objective 5.6.

Obstructed paper paths are often difficult to find. Usually it means disassembling the printer to find the bit of crumpled-up paper or other foreign substance that's blocking the paper path. A common obstruction is a piece of the "perf"—the perforated sides of tractor-feed paper—that has torn off and gotten crumpled up and then lodged in the paper path. It may be necessary to remove the platen roller and feed mechanism to get at the obstruction.

Stripped drive gears cause the paper to feed improperly, causing it to crinkle and cause jams. Using the wrong paper, such as thick paper when the platen has been set for thin paper, can also cause jams. When loading new paper, always be sure that the platen is properly adjusted.

Impact printers are used for multipart (or multipage) forms. Those forms are typically three or more sheets of paper thick. If the multipage forms are not feeding properly, it could be that the printer is set to receive paper that is too thin or too thick. Check the platen and adjust accordingly.

Use extra caution when printing peel-off labels in impact printers. If a label or even a whole sheet of labels becomes misaligned or jammed, do *not* roll the roller backward to realign the sheet. The small plastic paper guide that most impact printers use to control the forward movement of the paper through the printer will peel the label right off its backing if you reverse the direction of the paper. Once the label is free, it can easily get stuck under the platen, causing paper jams. A label stuck under the platen is almost impossible to remove without disassembling the paper-feed assembly. If a label is misaligned, try realigning the whole sheet of labels *slowly* using the feed roller (with the power off), moving it in very small increments.

Impact Printer Stepper Motor Problems

Printers use stepper motors to move the printhead back and forth as well as to advance the paper. The carriage motor is responsible for the back-and-forth motion while the main motor advances the paper. These motors get damaged when they are forced in any direction while the power is on. This includes moving the printhead over to install a printer ribbon as well as moving the paper-feed roller to align paper. These motors are very sensitive to stray voltages. If you are rotating one of these motors by hand, you are essentially turning it into a small generator and thus damaging it.

A damaged stepper motor is easy to detect. Damage to the stepper motor will cause it to lose precision and move farther with each step. If the main motor is damaged (which is more likely to happen), lines of print will be unevenly spaced. If the printhead motor goes bad, characters will be scrunched together. If a stepper motor is damaged badly enough, it won't move at all in any direction; it may even make grinding or high-pitched squealing noises. If any of these symptoms appear, it's time to replace one of these motors.

Stepper motors are usually expensive to replace—about half the cost of a new printer! Damage to them is easy to avoid; the biggest key is to not force them to move when the power is On.

Grinding noise, which is listed in Exam Objective 5.6, is a printer symptom that could occur across all printer types. Remember that to make a noise, a part has to move. So, think about motors, gears, and rollers if you hear grinding. It will probably help to get your ear close to the printer and/or remove the outer casing to see if you can tell exactly where it's coming from.

Inkjet Printer Problems

An inkjet printer has many of the same types of parts as an impact printer. In this sense, it's almost as though the inkjet technology is simply an extension of the technology used in impact printers. The parts on an inkjet can be divided into the following four categories:

- Printhead/ink cartridge
- Printhead carriage, belt, and stepper motor
- Paper-feed mechanism
- Control, interface, and power circuitry

Perhaps the most obvious difference between inkjet and impact printers is that impact printers often use tractor-feed paper, whereas inkjets use normal paper. The differences don't end there, though. Inkjet printers work by spraying ink (often in the form of a bubble) onto a page. The pattern of the bubbles forms images on the paper.

Inkjet printers are the most common type of printer found in homes because they are inexpensive and produce good-quality images. For this reason, you need to understand the most common problems with these printers so that your company can service them effectively. Let's take a look at some of the most common problems with inkjet printers and their solutions.

Print Quality

The majority of inkjet printer problems are quality problems. Ninety-nine percent of these can be traced to a faulty ink cartridge. With most inkjet printers, the ink cartridge contains the printhead and the ink. The major problem with this assembly can be described by "If you don't use it, you lose it." The ink will dry out in the small nozzles and block them if they are not used at least once every week or two.

Remember that three print quality issues, lines down the printed page, garbled print, and faded prints, are listed in Exam Objective 5.6.

An example of a quality problem is when you have thin, blank lines present in every line of text on the page. This is caused by a plugged hole in at least one of the small, pinhole ink nozzles in the print cartridge. Another common problem is faded printing. Replacing the ink cartridge generally solves these issues.

As we warned in Volume 1 Chapter 5, "Printers and Multifunction Devices," some people try to save money by refilling their ink cartridge when they need to replace it. If you are one of them, stop! Don't refill your ink cartridges! Almost all ink cartridges are *not* designed to be refilled. They are designed to be used once and disposed of (preferably sent to a recycling center). By refilling them, you make a hole in them. Ink can leak out, and the printer will need to be cleaned. The ink will probably also be of the wrong type, and print quality can suffer. Finally, using a refilled cartridge may void the printer's warranty.

If an ink cartridge becomes damaged or develops a hole, it can put too much ink on the page and the letters will smear. Again, the solution is to replace the ink cartridge. (You should be aware, however, that a very small amount of smearing is normal if the pages are laid on top of each other immediately after printing.)

One final print quality problem that does not directly involve the ink cartridge occurs when the print quickly goes from dark to light and then prints nothing. As previously mentioned, ink cartridges dry out if not used. That's why the manufacturers include a small suction pump inside the printer that primes the ink cartridge before each print cycle. If this priming pump is broken or malfunctioning, this problem will manifest itself and the pump will need to be replaced.

If the problem of the ink quickly going from dark to light and then disappearing ever happens to you and you really need to print a couple of pages, try this trick. First, take the ink cartridge out of the printer. Then squirt some window cleaner on a paper towel and gently tap the printhead against the wet paper towel. The force of the tap plus the solvents in the window cleaner should dislodge any dried ink, and the ink will flow freely again. Just be careful to not rub the paper towel across the printhead, because this could damage the nozzles.

After you install a new cartridge into many inkjet printers, the printheads in that cartridge must be aligned. *Printhead alignment* is the process by which the printhead is calibrated for use. A special utility that comes with the printer software is used to do this. Sometimes it's run from the printer itself and other times from the computer the printer is installed on. They vary a bit in how they work. For example, one utility might have the printer print several

vertical and horizontal lines with numbers next to them. The utility then displays a screen and asks you to choose the horizontal and vertical lines that are the most "in line." Once you enter the numbers, the software understands whether the printhead(s) are out of alignment, which direction, and by how much. The software then makes slight modifications to the print driver software to tell it how much to offset when printing. Other calibration software will print a pattern and then ask you to put the newly printed page on the scanner portion of the printer. It will then scan the pattern to make sure that the heads are properly aligned. Occasionally, alignment must be done several times to get the images to align properly.

Most new inkjet printers automatically align the printhead, and no interaction is required on your part. Even if this is the case, your printer software may have an option for you to be able to align the printheads manually.

Color Output Problems

Sometimes when you print a color document, the colors might not be the same colors that you expected based on what you saw on the screen. This is called an incorrect chroma display. A few different issues could cause this problem. First, ink could be bleeding from adjacent areas of the picture, causing the color to be off. A leaking cartridge can cause this, as can using the wrong type of paper for your printer.

If you know that you're using the right paper, try cleaning the print cartridges using the software utility that should have been included with the printer software. Once you do that, print a test page to confirm that the colors are correct. On most color printers, the test page will print colors in a pattern from left to right that mirrors the way the ink cartridges are installed. That brings us to our second potential problem: the ink cartridges are installed in the wrong spot. (This is for printers with multiple color ink cartridges.) That should be easy to check. Obviously, if that's the problem, put the color cartridges where they're supposed to be.

Third, if the ink that comes out of the cartridge doesn't match the label on the cartridge, try the self-cleaning utility. If that doesn't help, replace the cartridge. Finally, if one of the colors doesn't come out at all and self-cleaning doesn't help, just replace the cartridge.

Somewhat related to color problems is speckling on printed pages. This is where the pages have random dots of ink or other material on them as they print. This is most often caused by stuff like paper dust or residue from envelopes, staples, or glue getting into the machinery. Cleaning the printer and blowing it out with compressed air should clear up any speckling. It's possible, but less likely, that speckles could be caused by a leaky ink cartridge.

Speckling on printed pages is a symptom listed in Exam Objective 5.6.

Paper Jams

Inkjet printers have pretty simple paper paths, so paper jams due to obstructions are less likely than they are on impact printers. They are still possible, however, so an obstruction shouldn't be overlooked as a possible cause of jamming.

Paper jams in inkjet printers are usually due to one of two things:

- A worn pickup roller
- The wrong type of paper

The pickup roller usually has one or two D-shaped rollers mounted on a rotating shaft. When the shaft rotates, one edge of the D roller rubs against the paper, pushing it into the printer. When the roller gets worn, it gets smooth and doesn't exert enough friction against the paper to push it into the printer.

If the paper used in the printer is too smooth, it can cause the same problem. Pickup rollers use friction, and smooth paper doesn't offer much friction. If the paper is too rough, on the other hand, it acts like sandpaper on the rollers, wearing them smooth. Here's a rule of thumb for paper smoothness: Paper slightly smoother than a new dollar bill will work fine.

Creased paper is a common culprit in paper jams. The printer can crease the paper if there are obstructions in the paper path or problems with the paper-feed mechanism.

Paper-Feeding Problems

You will normally see one of two paper-feeding options on an inkjet printer. The first is that the paper is stored in a paper tray on the front of the printer. The second, which is more common on smaller and cheaper models, is for the paper to be fed in vertically from the back of the printer in a paper feeder. Both types may also have manual feed or envelope feed options.

Remember that paper jams, paper not feeding, and multipage misfeed are three issues listed in Exam Objective 5.6.

Regardless of the feed style, the printer will have a paper-feed mechanism, which picks up the paper and feeds it into the printer. Inside the paper-feed mechanism are pickup rollers, which are small rubber rollers that rub up against the paper and feed it into the printer. They press up against small rubber or cork patches known as *separation pads*. These help to keep the rest of the paper in the tray so that only one sheet is picked up at a time. A pickup stepper motor turns the pickup rollers.

If your printer fails to pick up paper, it could indicate that the pickup rollers are too worn. If your printer is always picking up multiple sheets of paper, it could be a couple of things, such as problems with the separation pads or your paper being too "sticky," damp,

or rough. Some printers that use vertical paper feeders have a lever with which you can adjust the amount of tension between the pickup rollers and the separation pads. If your printer is consistently pulling multiple sheets of paper, you might want to try to increase the tension using this lever.

The final component is the paper-feed sensor. This sensor is designed to tell the printer when it's out of paper, and it rarely fails. When it does, the printer will refuse to print because it thinks it is out of paper. Cleaning the sensor might help, but if not, you should replace the printer.

Inkjet Printer Stepper Motor Problems

Inkjet printers use stepper motors, just like impact printers. On an inkjet, the printhead carriage is the component containing the printhead that moves back and forth. A carriage stepper motor and an attached belt (the carriage belt) are responsible for the movement. The printhead carriage stays horizontally stable by resting on a metal stabilizer bar. Another stepper motor is responsible for advancing the paper.

Stepper motor problems on an inkjet printer will look similar to the ones on an impact printer. That is, if the main motor is damaged, lines of print will be unevenly spaced, and if the printhead motor goes bad, characters will be scrunched together. A lot of damage may cause the stepper motor to not move at all and possibly make grinding or high-pitched squealing noises. If any of these symptoms appear, it's time to replace one of these motors. As with impact printers, stepper motors can be expensive. It may make more economical sense to replace the printer.

Remember that grinding noises, which is an issue listed in Exam Objective 5.6, are often associated with motors on printers.

Power Problems

Inkjet printers have internal power circuits that convert the electricity from the outlet into voltages that the printer can use—typically, 12V and 5V. The specific device that does this is called the *transformer*. If the transformer fails, the printer will not power up. If this happens, it's time to get a new printer.

Laser Printer Problems

The process that laser printers use to print, called the *electrophotographic (EP) imaging process*, is the most complex process of all commonly used printers. We introduced this process in Volume 1 Chapter 4. It used to be specifically on the A+ Exam Objectives but was removed for 220-1201. You should still have familiarity with it, though, because knowing how the process works will help you figure out where or why it's not working. Table 13.2 provides a short description of what happens in each step.

TABLE 13.2 The EP imaging process

Step	Action
Processing	The page to be printed gets rendered, one horizontal strip at a time. The image is stored in memory for printing.
Charging	The charging corona gets a high voltage from the high-voltage power supply (HVPS). It uses the voltage to apply a strong uniform negative charge (–600VDC) to the photosensitive drum.
Exposing	The laser scans the drum. Wherever it touches the drum, the charge is reduced from –600VDC to around –100VDC. The pattern formed on the drum will be the image that is printed.
Developing	The developing roller acquires a –600VDC charge from the HVPS and picks up toner, which gets the same –600VDC charge. As the developing toner rolls by the photosensitive drum, the toner is attracted to the lesser-charged (–100VDC) areas on the photosensitive drum and sticks to it in those areas.
Transferring	The charging corona wire or roller acquires a strong positive charge (+600VDC) and transfers it to the paper. As the photosensitive drum with ink on it rolls by, the ink is attracted to the paper.
Fusing	The 350°F fuser roller melts the toner paper and the rubberized pressure roller presses the melted toner into the paper, making the image permanent.
Cleaning	A rubber blade scrapes any remaining toner off the drum, and a fluorescent lamp discharges any remaining charge on the photosensitive drum.

Looking at the steps involved in laser printing, it's pretty easy to tell that laser printers are the most complex printers that we have discussed. The good news, though, is that most laser printer problems are easily identifiable and have specific fixes. Let's discuss the most common laser and page printer problems and their solutions.

The Importance of Maintenance Kits

Don't forget to perform periodic preventive maintenance on your laser printers. Doing so can help eliminate many potential problems before they happen, and also improve print quality, reduce paper jams, and extend the life of the printer. Preventive maintenance includes cleaning the printer and using manufacturer-recommended maintenance kits.

Maintenance kits are a collection of replacement parts that are most likely to wear out, such as various rollers and pads, as well as the fuser assembly. If you're running into print quality issues or frequent paper jams, a maintenance kit is a great solution to try before going to more extreme measures such as replacing the printer.

Power Problems

If you turn on your laser printer and it doesn't respond normally, there could be a problem with the power it's receiving. Of course, the first thing to do is to ensure that it's plugged in.

A laser printer's DC power supply provides three different DC voltages to printer components. These can be checked at a power interface labeled J210, which is a 20-pin female interface. Pin 1 will be in the lower-left corner, and the pins along the bottom will all be odd numbers, increasing from left to right.

You can use a multimeter to test printer voltages.

Using the multimeter, you should find the following voltages:

- Pin 1 +5V
- Pin 5 –5V
- Pin 9 +24V

If none of the voltages are reading properly, then you probably need to replace the fuse in the DC power supply. If one or more (but not all) of the voltages aren't reading properly, then the first thing to do is to remove all optional hardware in the printer (including memory) and test again. If the readings are still bad, then you likely need to replace the DC power supply.

No Connectivity (IP Issues)

You can connect many laser printers directly to your network by using an Ethernet network cable (such as Category 5e, 6, or 6a) or by using a wireless network adapter with the printer. In cases like these, the printer acts as its own print server (typically, print server software is built into the printer), and it can speed up printing because you don't have a separate print server translating and then sending the directions to the printer.

Exam Objective 5.6 lists connectivity issues as a symptom to know. It can happen with inkjet and impact printers as well, but is most common with laser printers.

For printers such as these, no connectivity can be a sign of improperly configured IP settings, such as the IP address. While each printer is somewhat different, you can manually configure most laser printers' IP settings a number of ways:

- Through the printer's LCD control panel. For example, on several HP LaserJet models, you press Menu, navigate to the Network Config menu, select TCP/IP Config, select Manual, and then enter the IP address. You would then also configure the subnet mask and default gateway.

- By using Telnet to connect to the printer's management software from your computer. Note that Telnet is not secure. If you need a secure connection, use Secure Shell (SSH) instead.
- By using the management software that came with your printer.

You can also configure most IP printers to obtain an IP address automatically from a Dynamic Host Configuration Protocol (DHCP) server. When the printer is powered up, it will contact the server to get its IP configuration information, just like any other client on the network. While this may be convenient, it's usually not a good idea to assign dynamic IP addresses to printers. Client computers will have their printer mapped to a specific IP address; if that address is changed, you will have a lot of people complaining about no connectivity. If you are using the DHCP server to manage all of your network's IP addresses, be sure to reserve a static address for the printers. And, because printers may be moved from place to place in office environments, the device might end up on a different subnet with no connectivity.

To see the setting of a printer's IP address, print a configuration page from the printer's control panel. Then post the IP information near the printer so that users can easily connect to it.

Nothing Prints

You tell your computer to print, but nothing comes out of the printer. This problem is probably the most challenging to solve because several different things could cause it. Are you the only one affected by the problem, or are others having the same issue? Is the printer plugged in, powered on, and online? As with any troubleshooting, check your connections first.

Sometimes when nothing prints, you get a clue as to what the problem is. The printer may give you an "out of memory" error or something similar. Another possibility is that the printer will say "processing data" (or something similar) on its LCD display and nothing will print. It's likely that the printer has run out of memory while trying to process the print job. If your printer is exhibiting these symptoms, it's best to power the printer off and then power it back on.

Be aware that large print jobs may cause the printer to say "processing data" for several minutes before the print job starts. There is nothing wrong with this, although it's possible that your printer could stand a memory upgrade. But if the printer exhibits this behavior for a long time, say 20 or 30 minutes, it may be best to cycle the power.

Paper Jams

Laser printers today run at copier speeds. Because of this, their most common problem is paper jams. Paper can get jammed in a printer for several reasons. First, feed jams happen when the paper-feed rollers get worn (similar to feed jams in inkjet printers). The solution to this problem is easy: replace the worn rollers.

Remember that paper jams, paper not feeding, and multipage misfeed are issues listed in Exam Objective 5.6.

Another cause of feed jams is related to the drive gear of the pickup roller. The drive gear (or clutch) may be broken or have teeth missing. Again, the solution is to replace it. To determine if the problem is a broken gear or worn rollers, print a test page, but leave the paper tray out. Look into the paper-feed opening with a flashlight and see if the paper pickup roller(s) are turning evenly and don't skip. If they turn evenly, the problem is probably worn rollers.

If worn pickup rollers are causing your paper-feed jams, there is something that you can do to get your printer working while you're waiting for the replacement pickup rollers or maintenance kit. Scuff the feed rollers with a steel wool pot-scrubber pad (or something similar) to roughen them up. This trick works only once. After that, the rollers aren't thick enough to touch the paper.

Worn exit rollers can also cause paper jams. These rollers guide the paper out of the printer into the paper-receiving tray. If they are worn or damaged, the paper may catch on its way out of the printer. These types of jams are characterized by a paper jam that occurs just as the paper is getting to the exit rollers. If the paper jams, open the rear door and see where the paper is located. If the paper is very close to the exit rollers, they are probably the problem.

The solution is to replace all the exit rollers. You must replace them all at the same time because even one worn exit roller can cause the paper to jam. Besides, they're inexpensive. Don't skimp on these parts if you need to have them replaced.

Paper jams can also be the fault of the paper. If your printer consistently tries to feed multiple pages into the printer, the paper isn't dry enough. If you live in an area with high humidity, this could be a problem. We've heard some solutions that are pretty far out but work (like keeping the paper in a Tupperware-type airtight container or microwaving it to remove moisture). The best all-around solution, however, is humidity control and keeping the paper wrapped until it's needed. Keep the humidity around 50 percent or lower (but above 25 percent if you can, in order to avoid problems with electrostatic discharge).

Real World Scenario

Printer Triage

One of the authors relates the following story. He was in the local hospital ER a while ago having his hand examined. The receptionist asked him a few questions, filled out a report in the medical database on her computer, and printed it. When the paper starting coming out of the laser printer, she grabbed it and "ripped" it from the printer, as you might do if the paper were in an old typewriter. The printer's exit rollers complained bitterly and made a noise that made him cringe. She did this for every sheet of paper she printed.

The following week, that printer came in for service because it was jamming repeatedly. The problem? Worn exit rollers.

He had a word with the person in charge of computer repair at that hospital and saved them from many future repairs. The lesson? Printers don't have to be treated with kid gloves, but using them properly can prolong their life and reduce the need for service repairs.

Finally, a grounded metal strip called the *static-charge eliminator strip* inside the printer drains the transfer corona charge away from the paper after it has been used to transfer toner from the EP cartridge. If that strip is missing, broken, or damaged, the charge will remain on the paper and may cause it to stick to the EP cartridge, causing a jam. If the paper jams after reaching the transfer corona assembly, this may be the cause.

Blank Pages

It's really annoying to print a 10-page contract and receive 10 pages of blank paper from the printer. Blank pages are a somewhat common occurrence in laser printers. Somehow, the toner isn't being put on the paper. There are three major causes of blank pages:

Toner Cartridge The toner cartridge is the source of most quality problems because it contains most of the image-formation pieces for laser printers. Let's start with the obvious. A blank page or faded prints will come out of the printer if there is no toner or low toner in the cartridge. It might sound simple, but some people think these things last forever. Many laser printers give some sort of warning if the toner cartridge is low, but it's easy to check. Just open the printer, remove the toner cartridge, and shake it. You will be able to hear if there's toner inside the cartridge. If it's empty, replace it with a known-good, manufacturer-recommended toner cartridge. If it is not yet empty, shaking it redistributes the toner and may provide better printing for some time.

When you're shaking a toner cartridge, loose toner can fall out of the cartridge and get on your clothing. Always hold the toner cartridge away from your body when shaking it.

Another issue that crops up rather often is the problem of using refilled or reconditioned toner cartridges. During their recycling process, these cartridges may be filled with the wrong kind of toner (for example, one with an incorrect composition). This can cause toner to be repelled from the EP drum instead of being attracted to it. Thus, there's no toner on the page because there was no toner on the EP drum to begin with. The solution once again is to replace the toner cartridge with the type recommended by the manufacturer.

A third problem related to toner cartridges happens when someone installs a new toner cartridge and forgets to remove the sealing tape that is present to keep the toner in the cartridge during shipping. The solution to this problem is as easy as it is obvious: remove the toner cartridge from the printer, remove the sealing tape, and reinstall the cartridge.

Most of the time, if you have dust or debris in a printer, you can use compressed air to blow it away. Don't do that with toner, though, because it will make a huge mess. If you have a toner spill, use a specialized toner vacuum to pick it up. Also, never use a damp cloth to try to clean up a toner spill. If a cloth is needed, use a dry one.

Transfer Corona Assembly The second cause of the blank-page problem is a damaged or missing transfer corona wire or damaged transfer corona roller. If a wire is lost or damaged, the developed image won't transfer from the EP drum to the paper; thus, no image appears on the printout. To determine if this is causing your problem, do a self-test (you may need to look up how to do that for your printer). If there is an image on the drum but not on the paper, you know that the transfer corona assembly isn't doing its job.

To check if the transfer corona assembly is causing the problem, open the cover and examine the wire (or roller, if your printer uses one). The corona wire is hard to see, so you may need a flashlight. You will know if it's broken or missing just by looking at it. (It will either be in pieces or just not be there.) If it's not broken or missing, the problem may be related to the high-voltage power supply.

The transfer corona wire (or roller) is a relatively inexpensive part. You can easily replace it by removing two screws and having some patience.

High-Voltage Power Supply The high-voltage power supply (HVPS) supplies high-voltage, low-current power to both the charging and transfer corona assemblies in laser printers. If it's broken, neither corona will work properly. If a self-test shows an image on the drum but none on the paper, and the transfer corona assembly is present and not damaged, then the HVPS is at fault.

All-Black Pages

Only slightly more annoying than 10 blank pages are 10 black pages. This happens when the charging unit (the charging corona wire or charging corona roller) in the toner cartridge malfunctions and fails to place a charge on the EP drum. Because the drum is grounded, it has no charge. Anything with a charge (like toner) will stick to it. As the drum rotates, all of the toner is transferred to the page and a black page is formed.

This problem wastes quite a bit of toner, but it can be fixed easily. The solution (again) is to replace the toner cartridge with a known, good, manufacturer-recommended one. If that doesn't solve the problem, then the HVPS is at fault. (It's not providing the high voltage that the charging corona needs to function.)

Repetitive Small Marks or Defects

Repetitive marks occur frequently in heavily used (as well as older) laser printers. Toner spilled inside the printer may be causing the problem. It can also be caused by a crack or chip in the EP drum (this mainly happens with recycled cartridges), which can accumulate toner. In both cases, some of the toner gets stuck onto one of the rollers. Once this happens, every time the roller rotates and touches a piece of paper, it leaves toner smudges spaced a roller circumference apart.

The solution is relatively simple: Clean or replace the offending roller (such as by using a maintenance kit). To help you figure out which roller is causing the problem, the service manuals contain a chart like the one shown in Figure 13.14. (Some larger printers also have the roller layout printed inside the service door.) To use the chart, place the printed page next to it. Align the first occurrence of the smudge with the top arrow. The next smudge will line up with one of the other arrows. The arrow it lines up with tells you which roller is causing the problem.

FIGURE 13.14 Laser printer roller circumference chart

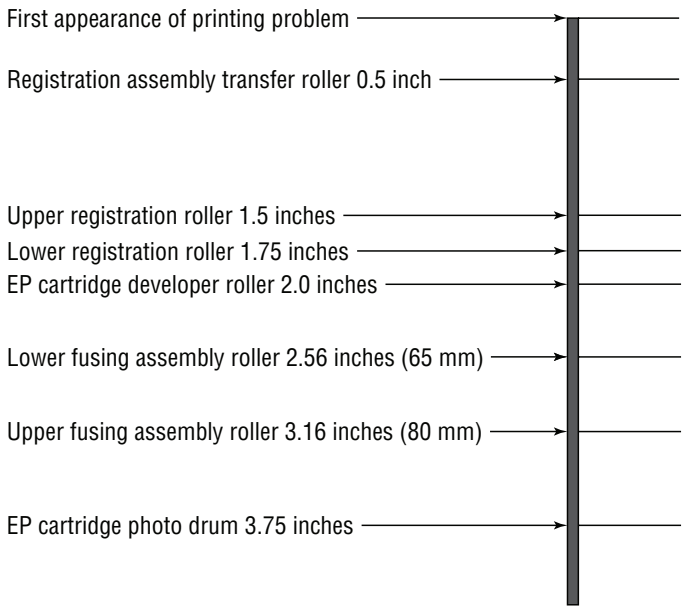
First appearance of printing problem
Registration assembly transfer roller 0.5 inch
Upper registration roller 1.5 inches
Lower registration roller 1.75 inches
EP cartridge developer roller 2.0 inches
Lower fusing assembly roller 2.56 inches (65 mm)
Upper fusing assembly roller 3.16 inches (80 mm)
EP cartridge photo drum 3.75 inches

Remember that the chart in Figure 13.14 is only an example. Your printer may have different-sized rollers and thus need a different chart. Check your printer's service documentation for a chart like this. It is valuable in determining which roller is causing a smudge.

Vertical White Lines on the Page

Vertical white lines running down all or part of the page are a relatively common problem on older printers, especially ones that don't see much maintenance. Foreign matter (more than likely toner) caught on the transfer corona wire causes this. The dirty spots keep the toner from being transmitted to the paper (at those locations, that is), with the result that streaks form as the paper progresses past the transfer corona wire.

The solution is to clean the corona wires. Many laser printers contain a small corona wire brush to help with this procedure. It's usually a small, green-handled brush located near the transfer corona wire. To use it, remove the toner cartridge and run the brush in the charging corona groove on top of the toner cartridge. Replace the cartridge, and then use the brush to remove any foreign deposits on the transfer corona. Be sure to put it back in its holder when you're finished.

Remember that lines down the printed pages is an issue listed in Exam Objective 5.6. It can be white lines or black lines.

Vertical Black Lines on the Page

A groove or scratch in the EP drum can cause the problem of vertical black lines running down all or part of the page. Because a scratch is lower than the surface, it doesn't receive as much (if any) of a charge as the other areas. The result is that toner sticks to it as though it were discharged. The groove may go around the circumference of the drum, so the line may go all the way down the page.

Another possible cause of vertical black lines is a dirty charging corona wire. A dirty charging corona wire prevents a sufficient charge from being placed on the EP drum. Because the charge on the EP drum is almost zero, toner sticks to the areas that correspond to the dirty areas on the charging corona.

The solution to the first problem is, as always, to replace the toner cartridge (or EP drum, if your printer uses a separate EP drum and toner). You can also solve the second problem with a new toner cartridge, although that would be an extreme solution. It's easier to clean the charging corona with the brush supplied with the cartridge.

Image Smudging

If you can pick up a sheet from a laser printer, run your thumb across it, and have the image come off on your thumb, then you have a fuser problem. The fuser isn't heating the toner and fusing it into the paper. This could be caused by a number of things—but all of them can

be handled by a fuser replacement. For example, if the halogen light inside the heating roller has burned out, that would cause the problem. The solution is to replace the fuser. The fuser can be replaced with a rebuilt unit, if you prefer. Rebuilt fusers are almost as good as new ones, and some even come with guarantees. Plus, they cost less.

The whole fuser may not need to be replaced. Fuser components can be ordered from parts suppliers and can be rebuilt by you. For example, if the fuser has a bad lamp, you can order a lamp and replace it in the fuser.

A similar problem occurs when small areas of smudging repeat themselves down the page. Dents or cold spots in the fuser heat roller cause this problem. The only solution is to replace either the fuser assembly or the heat roller.

Ghosting

Ghosting (or echo images) is what you have when you can see faint images of previously printed pages on the current page. This is caused by one of two things: a broken cleaning blade or bad erasure lamps. A broken cleaning blade causes old toner to build up on the EP drum and consequently present itself in the next printed image. If the erasure lamps are bad, then the previous electrostatic discharges aren't completely wiped away. When the EP drum rotates toward the developing roller, some toner sticks to the slightly discharged areas.

Double/echo images on the print is an issue listed in Exam Objective 5.6.

If the problem is caused by a broken cleaner blade, you can replace the toner cartridge. If it's caused by bad erasure lamps, you'll need to replace them. Because the toner cartridge is the least expensive cure, you should try that first. Usually, replacing the toner cartridge will solve the ghosting problem. If it doesn't, you will have to replace the erasure lamps.

Printer Prints Pages of Garbage

This has happened to everyone at least once. You print a one-page letter, but instead of the letter you have 10 pages of what looks like garbage (or garbled characters) or many more pages with one character per page come out of the printer. This problem comes from one of two different sources:

Printer Driver The correct printer driver needs to be installed for the printer and operating system. For example, if you have an HP LaserJet Pro MFP M148 and a Windows computer, then you need to install an HP LaserJet Pro MFP M148 driver made for Windows. Once the driver has been installed, it must be configured for the correct page-description language: PCL or PostScript. Most HP LaserJet printers use PCL (but can be configured for PostScript). Determine which page-description language your printer has

been configured for, and set the printer driver to the same setting. If this is not done, you will get garbage out of the printer.

Most printers that have LCD displays will indicate that they are in PostScript mode with a *PS* or *PostScript* somewhere in the display.

If the problem is the wrong driver setting, the garbage that the printer prints will look like English. That is, the words will be readable, but they won't make any sense.

Formatter Board The other cause of several pages of garbage being printed is a bad formatter board. This circuit board turns the information the printer receives from the computer into commands for the various components in the printer. Usually, problems with the formatter board produce wavy lines of print or random patterns of dots on the page.

Replacing the formatter board in a laser printer is relatively easy. Usually, this board is installed under the printer and can be removed by loosening two screws and pulling it out. Typically, replacing the formatter board also replaces the printer interface, which is another possible source of garbage printouts.

Finishing Issues

Many laser printers today are multifunction devices that include copying and scanning. Some devices come with finishers, which add touches like collating, stapling papers together, or hole punching output. If one of those functions is not working properly, it's an issue with the finisher. More often than not, it's a simple cleaning that takes care of the issue. For example, if documents are not being stapled, it could be a jam in the stapling mechanism that needs to be cleared. Or, of course, the printer may be out of staples.

Exam Objective 5.6 lists finishing issues, including staple jams and hole punch, as areas you should be able to troubleshoot.

Managing Print Jobs

Most people know how to send a job to the printer. Clicking File and then Print, or pressing Ctrl+P on your keyboard, generally does the trick. But once the job gets sent to the printer, what do you do if it doesn't print?

Keep in mind that in a networked environment, users need the proper permissions both to install and to print to the printer. Not having permission will result in an "access denied" message (or something similar) and the inability to print.

When you send a job to the printer, that print job ends up in a line with all other documents sent to that printer. A series of print jobs waiting to use the printer is called the *print queue*. In most cases, the printer will print jobs on a first-come, first-served basis. (There are exceptions if you've enabled printing priorities on the Printer Properties Advanced tab.) Once you send the job to the printer in Windows, a small printer icon will appear in the notification area in the lower-right corner of your desktop, near the clock. By double-clicking it (or by right-clicking it and selecting the printer name), you will end up looking at the jobs in the print queue, like the one shown in Figure 13.15.

FIGURE 13.15 Print jobs in the print queue in Windows

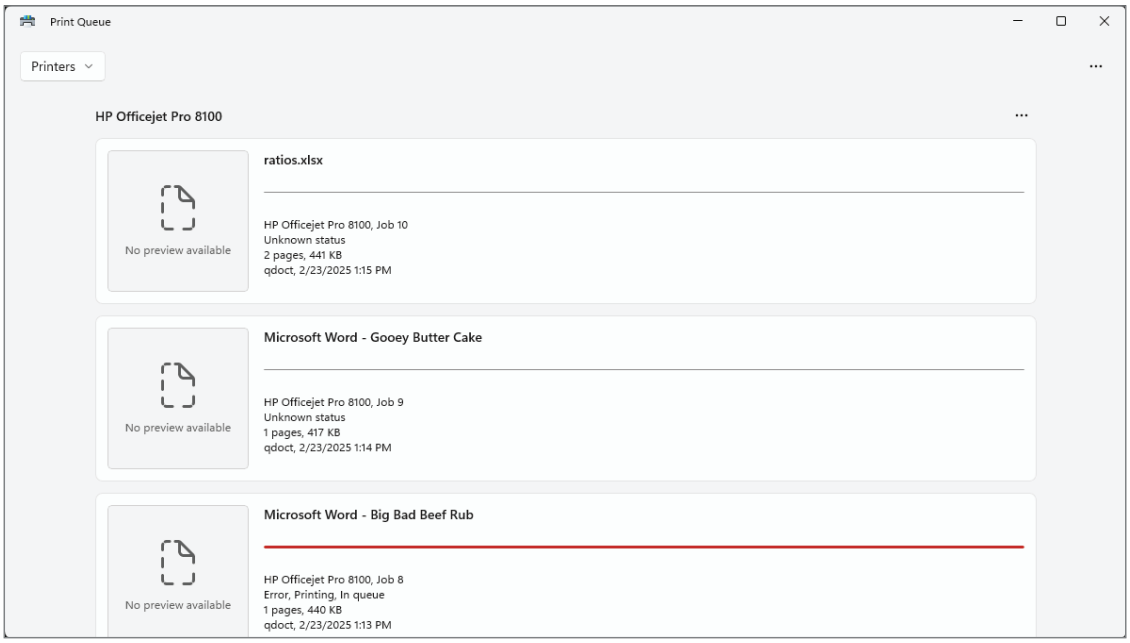

In Figure 13.15, you can see that the first document submitted (at the bottom of the list) has an error, which may explain why it hasn't printed. All the other documents in the queue are blocked until the job with the error is cleared. You can clear it one of two ways. Either right-click the document and choose Cancel (or Restart if you want to try printing it again), or click on the ellipsis (three dots) menu to the right of the printer name (Figure 13.16) and cancel all print jobs.

Once you have cleared the print job causing the problem, the next job will move to the top of the queue. It should show its status as Printing, like the one shown in Figure 13.17. But what if it shows that it's printing but it still isn't working? (We're assuming that the printer is powered on, connected properly, and online.) It could be a problem with the print spooler.

FIGURE 13.16 Printer ellipsis menu

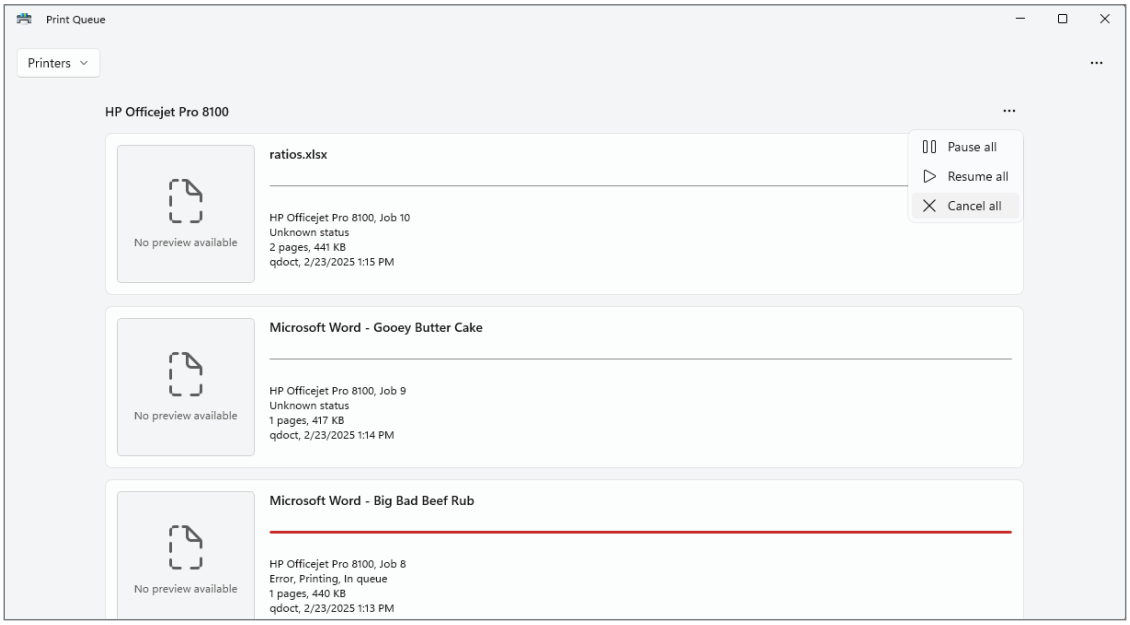

FIGURE 13.17 Print job printing correctly

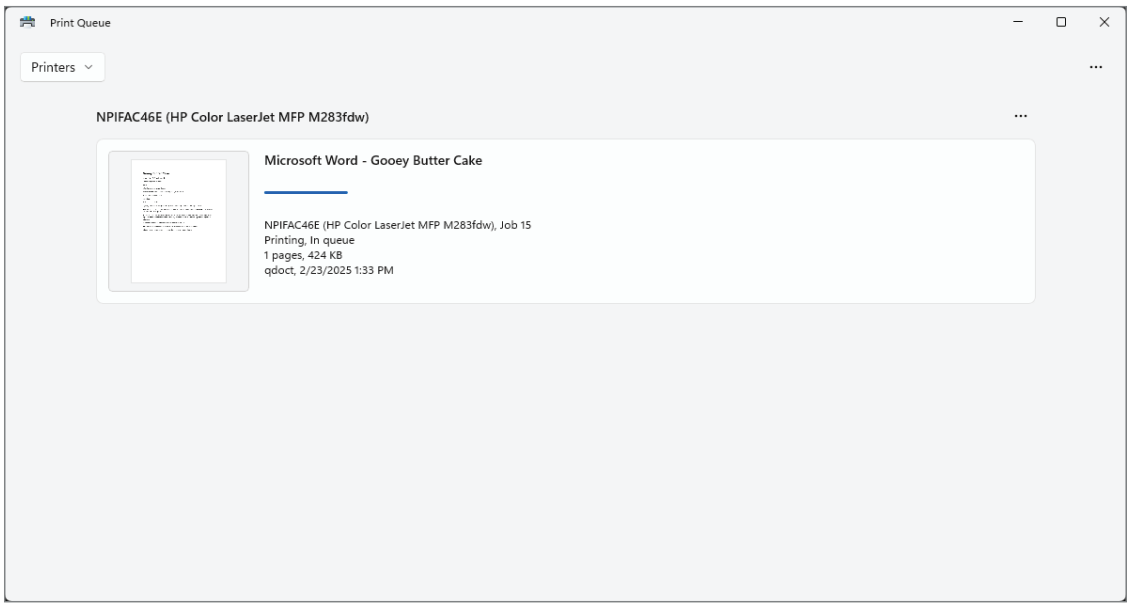
Print Queue
Printers
NPIFAC46E (HP Color LaserJet MFP M283fdw)
Microsoft Word - Gooey Butter Cake
NPIFAC46E (HP Color LaserJet MFP M283fdw), Job 15
Printing, In queue
1 pages, 424 KB
qdoct, 2/23/2025 1:33 PM

If print jobs are processed very slowly, or if you are continually seeing "low memory" error messages, it might be a good time to upgrade the memory in the printer.

Managing the Print Spooler

The *print spooler* is a service that formats print jobs in a language that the printer understands. Think of it as a holding area where the print jobs are prepared for the printer. In Windows, the spooler is started automatically when Windows loads.

Exam Objective 5.6 lists two issues related to print jobs: Multiple prints pending in queue and frozen print queue. Both are related to the print spooler.

If jobs aren't printing and there's no apparent reason why, it could be that the print spooler has stalled. To fix the problem, you need to stop and restart the print spooler. Exercise 13.1 walks you through stopping and restarting the spooler in Windows 11.

EXERCISE 13.1

Stopping and Restarting the Print Spooler in Windows 11

1. Open the Services app by typing **services** into the Windows 11 search bar and pressing Enter.
2. In the right pane, find the Print Spooler service, which is selected in Figure 13.18.

FIGURE 13.18 The Print Spooler service

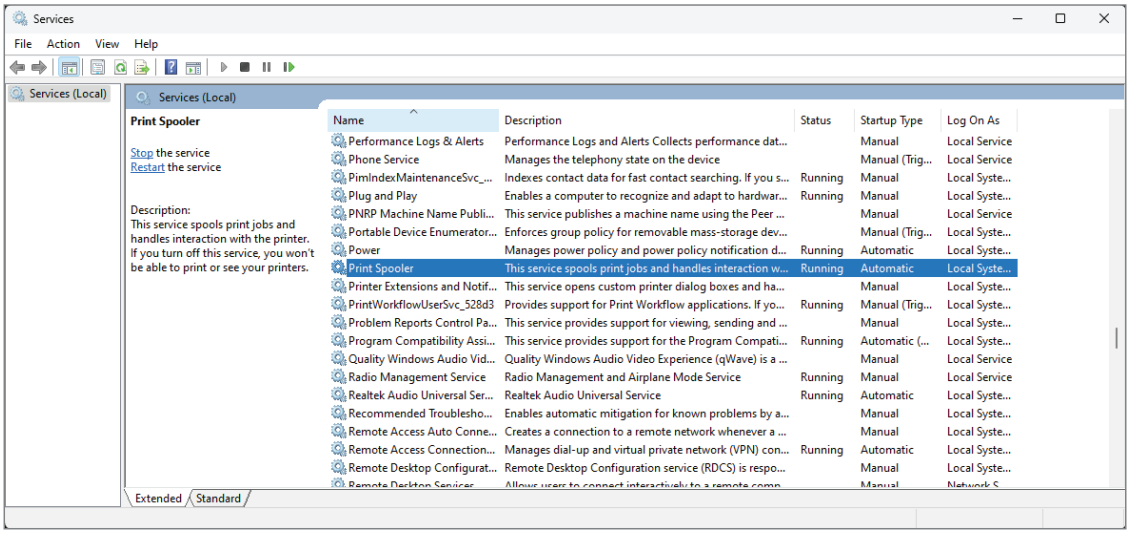

3. Stop the spooler. There are several ways to do this. You can right-click the service name and choose Stop, click the Stop square above the list of services, or click Stop under where it says Print Spooler to the left of the services list.

4. Restart the spooler by right-clicking the service name and choosing Start or by clicking the Start arrow above the list of services. After the spooler has restarted, the service's Status column should display Started.
5. Close the Services app.

If you have a different version of Windows, the steps to stop and restart the spooler are the same as in Exercise 13.1; the only difference might be in how you get to the Services app.

If you have multiple failed print jobs in system logs, such as in Windows Event Viewer, stopping and restarting the print spooler is a good solution to try. If the printer is wireless, the failed jobs could also be a sign of intermittent connectivity issues.

Printing a Test Page

If your printer isn't spitting out print jobs, it may be a good idea to print a test page to see if that works. The test page information is stored in the printer's memory, so there's no formatting or translating of jobs required. It's simply a test to make sure that your printer hears your computer.

When you install a printer, one of the last questions you're asked is whether to print a test page. If there's any question, go ahead and do it. If the printer is already installed, you can print a test page a few different ways. One is if you go to Settings ➢ Bluetooth & Devices ➢ Printers & Scanners and select the printer, there will be an option called Print Test Page. If you are in the printer's Properties window, there is a Print Test Page button on the General tab. If nothing happens, double-check your connections and stop and restart the print spooler. If garbage prints, there is likely a problem with the printer or the print driver.

Managing Print Options

Printers can print to paper of various sizes, as well as multiple orientations. If a print job is not coming out correctly, such as the output is squeezed into a smaller area than a regular sheet of 8.5 × 11 paper, it may be that the paper size is improperly set. Or, if the right side or bottom is cut off, it could be a page orientation issue. Both options are configured in the Printing preferences settings for the printer. To change the settings in Windows, open the Printers & Scanners app. Select the printer and then select Printing Preferences. It will open a window with the printer's name at the top, as shown in Figure 13.19.

In Figure 13.19, there is an option for paper sizes. (Note that this tab will look different for different printer models.) Letter paper is 8.5 × 11 inches, which is common in the United States. In most of Europe, the standard is called A4, which is 210 × 297 mm, or about 8.3 × 11.7 inches. Setting a printer to print to A4 when it should be Letter, or vice versa, can result in stretched or scrunched printing, or unusually larger than normal margins, depending on the printer. This tab also has settings for paper source. Some printers have multiple paper

trays. They can be configured such that one has letter size paper and another has legal (8.5 × 14 inches) size paper. When the user prints, they need to specify letter or legal, and the printer will pull paper from the correct tray.

FIGURE 13.19 Printing Preferences

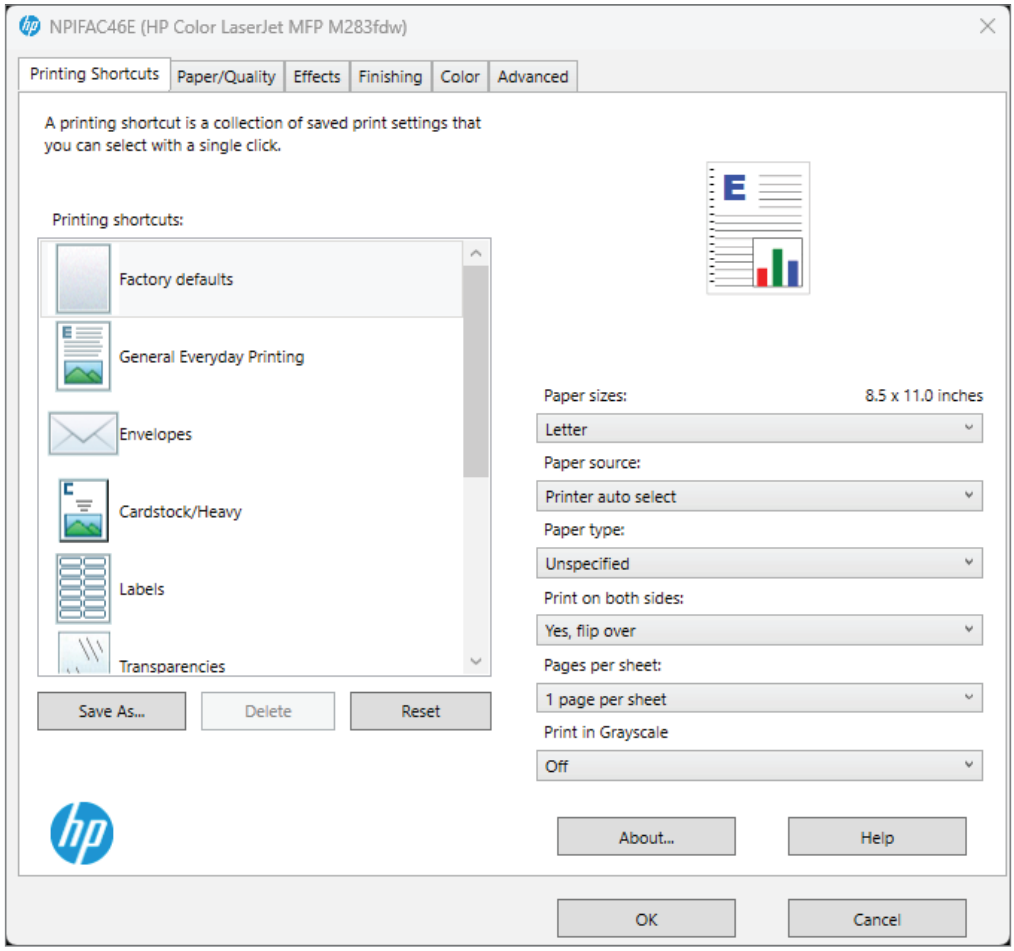

Incorrect page orientation and tray not recognized are two issues you should be able to troubleshoot, per Exam Objective 5.6.

If the paper tray has not been configured properly, either the printer won't recognize the tray (meaning that users won't be able to access it when they try to print) or it will pull paper from that tray thinking it's one size when it should be another. For example, if Tray 2 is loaded with letter paper, but the Printer Properties shows it as legal paper, users printing legal-sized documents using that tray will have the bottom three inches of their print jobs cut off. On this printer, paper trays are configured on the Paper/Quality tab, shown in Figure 13.20. You would select the Paper Source tray, and then configure the appropriate paper size for that tray using the Paper Sizes drop-down menu.

FIGURE 13.20 Paper/Quality tab

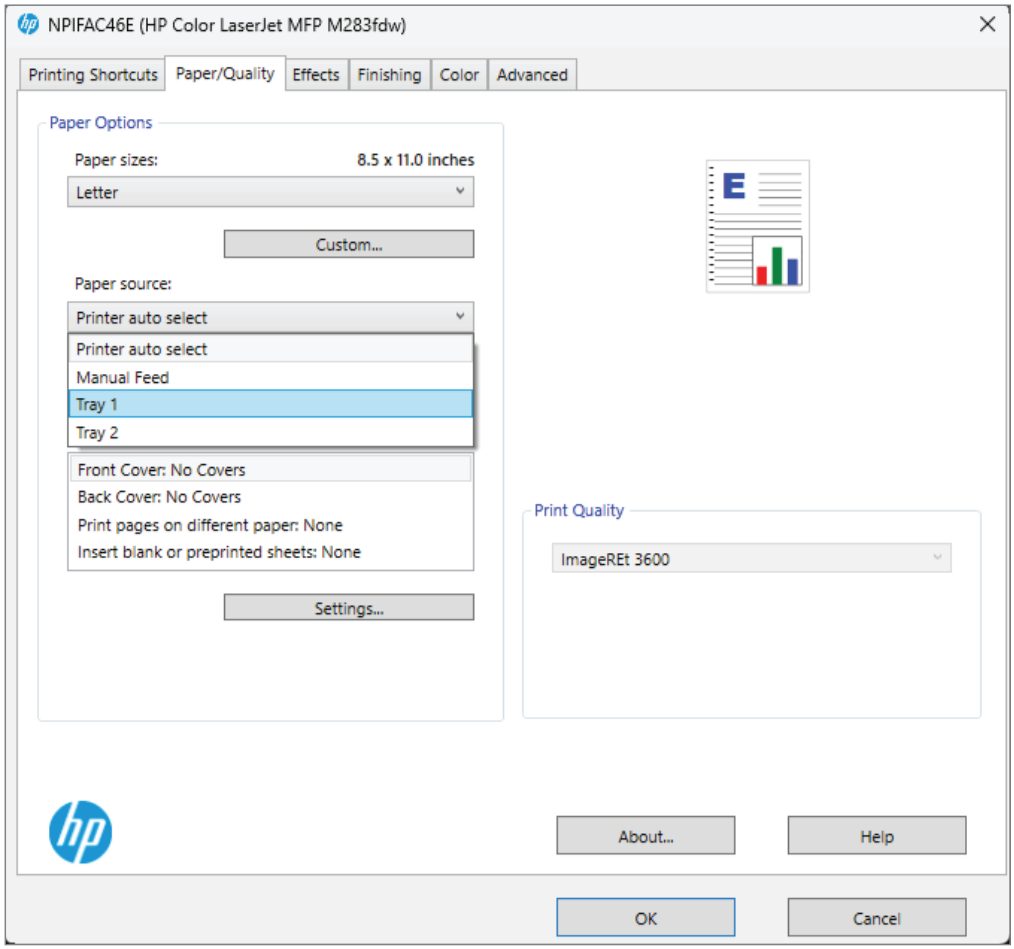

Page orientation is handled on the Finishing tab for this printer, shown in Figure 13.21. The two options are portrait (taller than it is wide) or landscape (wider than it is tall). Note as well that this printer has options for printing on both sides, a booklet layout, and number of pages per sheet. Again, different printers will have different options—you may just need to click on a few different tabs to find the setting you're looking for, if it's available on that printer.

Summary

In this chapter, we discussed mobile device, display, and printer troubleshooting. We started with video issues, such as input problems and image problems, and looked at how to resolve them. Next, we covered problems that are unique to mobile devices and laptop computers. Because of their compact nature, they have unique issues relating to heat and power, input and output, connectivity, and potential damage.

FIGURE 13.21 Finishing tab, including page orientation

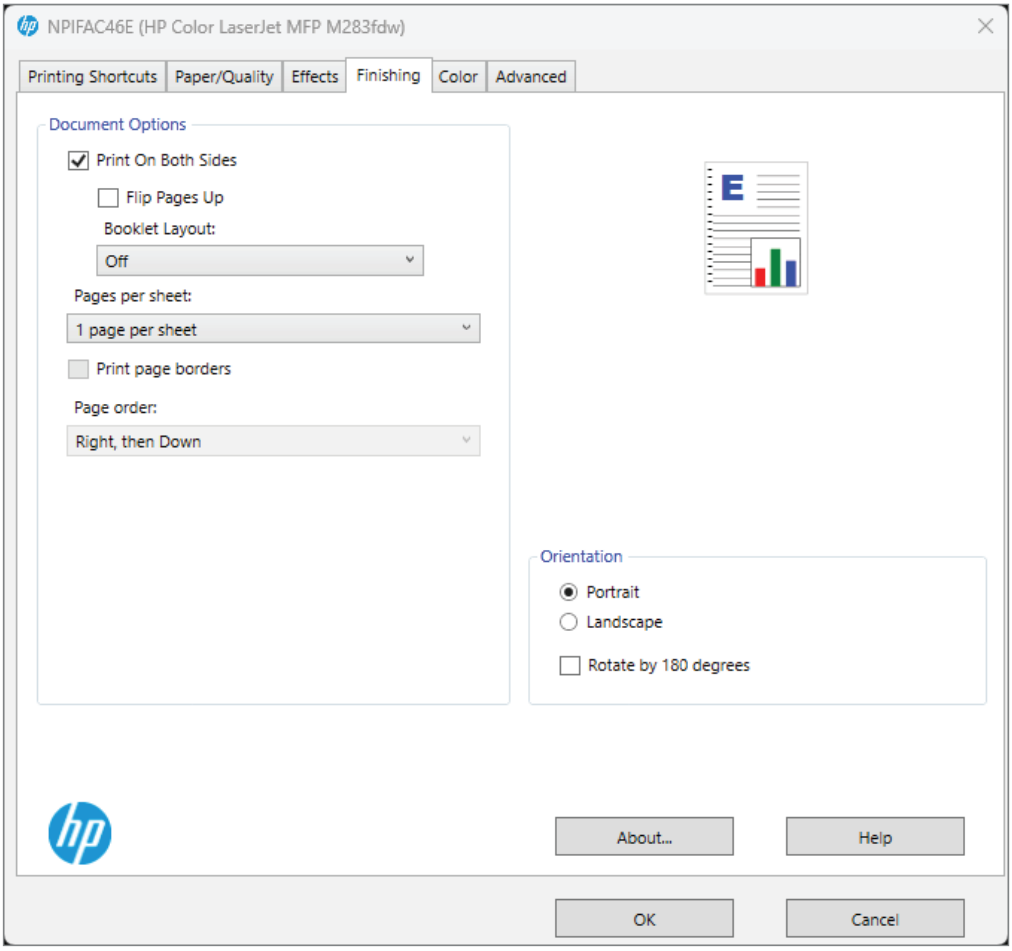

We followed that with a discussion on troubleshooting printers. Specifically, we explored problems with three major classes of printers—impact, inkjet, and laser—and then we talked about managing print jobs, the print spooler, printing a test page, and printer configuration options.

Exam Essentials

Understand common video, projector, and display input issues. Two input issues video displays might have are incorrect input source and physical cabling issues. If the source is incorrect, there will be no display. This is usually fixed by changing the input source from a menu on the display or within the operating system. If there is a physical cabling issue, the cable likely needs to be replaced.

Understand common video, projector, and display image problems. There are a variety of image problems a video display could have. They include a burnt-out bulb, fuzzy image, sizing issues, distorted image, display burn-in, dead pixels, flashing screen or flickering screen, incorrect color display, or dim image. These will either be fixed through a setting or by replacing defective hardware.

Understand other issues video, projector, and display systems may have. Many video systems have audio output, so audio issues may be related to the display device itself. Another issue is intermittent projector shutdown, which is typically related to overheating.

Understand common mobile device power and heat issues. Power issues can include a battery not charging, no power, extremely short battery life (poor battery health), overheating, and a swollen battery.

Know what to do to resolve common mobile device input problems. These can include sticking keys, ghost cursor or pointer drift, num lock lights, a stylus that doesn't work, or an unresponsive digitizer/touch screen. Be familiar with the laptop Function key and the features it controls, including switching the display to an external monitor.

Understand other issues common to mobile devices. Other issues to know include physical damage such as from dropping or liquid damage, broken screens, physically damaged ports, malware, unable to install new applications, and degraded performance.

Know how to set IP addresses on a printer. The IP address can often be obtained automatically from a DHCP server, but this is not recommended for corporate networks. Instead, use a static IP address (or use a DHCP reservation). You may be able to use the printer's control panel or printer management software to configure the IP address.

Know what could cause the printer to print garbage. Most often, the print driver causes this. Deleting and reinstalling it should fix the problem. A defective formatter board can also cause it.

Understand what could cause print-quality issues on an impact printer. Print-quality issues are generally related to either the ribbon or the printhead. The specific problem you are having will help determine the culprit.

Know what can cause unevenly spaced lines or characters on a dot-matrix or inkjet printer. A failing stepper motor usually causes this. For line spacing problems, it's the main stepper motor. For character spacing, it will be the carriage stepper motor.

Know what causes printers to have paper problems. On an impact printer, paper jams are usually caused by material getting into the rollers, such as extra "perf" from the tractor-feed paper. On inkjets and laser printers, worn pickup rollers often cause this problem. These same mechanisms can affect paper not feeding as well or multipage misfeeds.

Be familiar with troubleshooting a variety of printer output quality issues. Issues to know include lines down the printed pages, garbled prints, faded prints, speckling on printed pages, double/echo images on the print, finishing issues such as staple jams or hole punch problems, incorrect page orientation, or the printer tray not being recognized.

Understand print spooler issues and know how to stop and restart the print spooler. Print spooler issues can include multiple prints pending in the queue and a frozen print queue. To fix these, stop and restart the print spooler.

Open the Services applet. Find Print Spooler on the right side. Right-click it and click Stop, or highlight it and click the Stop square above the list of services. To restart the print spooler, right-click it and select Start, or click the Start triangle above the list of services.

Review Questions

The answers to the chapter review questions can be found in Appendix A.

1. If the video on your laptop is not working, what should you do to troubleshoot it? (Choose two.)
   - **A.** Toggle the video function key.
   - **B.** Try using an external monitor.
   - **C.** Remove the display unit and reattach it.
   - **D.** Power the system off and back on.

2. A printer is printing unusually large margins on documents that are sent to it. You suspect that the paper size is set incorrectly. Where should you go to change this?
   - **A.** Print spooler
   - **B.** Paper selection switch on the printer
   - **C.** Printing preferences
   - **D.** Print queue

3. A user tells you that his Android phone case is bulging, and the device has been hot to the touch lately. You suspect a swollen battery. Which of the following is a good next step?
   - **A.** Replace the device.
   - **B.** Leave the device on to drain the battery, and then recharge the device.
   - **C.** Turn the device off, let it cool to room temperature, and then recharge the device.
   - **D.** Use the battery calibration utility.
   - **E.** Freeze the system.

4. Users are complaining that their print jobs are not printing. You open the print queue and see 50 jobs lined up. The printer is connected properly and online. What should you do?
   - **A.** Open the Printer Troubleshooting utility on the printing device and have it diagnose the problem.
   - **B.** Stop and restart the print spooler.
   - **C.** Delete and reinstall the printer.
   - **D.** Delete and reinstall Windows.

5. The display on a laptop computer is continuously flickering. Which of the following are most likely to cause this problem? (Choose two.)
   - **A.** A failing display
   - **B.** Interference from a nearby fan
   - **C.** A failing backlight
   - **D.** A failing inverter
   - **E.** Incorrect video driver

6. You are having problems with the video card in one of your computers. Where should you check for troubleshooting information?
   - **A.** Another computer with the same video card
   - **B.** The video card manufacturer's website
   - **C.** The manual that came with the card
   - **D.** The server log

7. Your laser printer keeps printing vertical black lines on its output pages. What is the most likely cause of the problem?
   - **A.** There is a groove or scratch in the EP drum.
   - **B.** The EP drum-cleaning blade is broken.
   - **C.** The printer is low on toner.
   - **D.** The transfer corona wire is not working properly.

8. The display on your laptop appears warped and fuzzy. You plug in an external monitor, and the image on it is fine. What is the most likely cause of the problem?
   - **A.** The video card
   - **B.** The LCD display
   - **C.** The motherboard
   - **D.** The video driver

9. You have an inkjet printer. Recently, papers are being printed with excessive amounts of ink, and the ink is smearing. What is the most likely cause of the problem?
   - **A.** A faulty ink cartridge
   - **B.** A corrupted print driver
   - **C.** A faulty fuser
   - **D.** Too much humidity in the air

10. When you print documents on your laser printer, you see residue from previous images on the output. Which two things are the most likely causes of this problem? (Choose two.)
    - **A.** A faulty transfer corona wire
    - **B.** An overheating printer
    - **C.** A bad erasure lamp
    - **D.** A broken cleaning blade

11. You support an old impact printer at work. When the printer prints, there is always a blank horizontal line in the middle of each line of output. What is the most likely cause of the problem?
    - **A.** The print ribbon is old and needs to be replaced.
    - **B.** The print ribbon is not advancing properly.
    - **C.** The printhead needs to be replaced.
    - **D.** The wrong print driver is installed.

12. A user brings in an iPhone and the touchscreen is not responding. What should you do to try first to get the device functioning again?
    - **A.** Hold the power button and the sleep/wake button for 10 seconds to reset the device.
    - **B.** Restore the device to factory settings.
    - **C.** Replace the battery in the device.
    - **D.** Hold the power button for 10 seconds to force the device to power off.

13. The stylus for your manager's tablet has stopped working. The tablet responds to input when you touch the screen with your finger. Which troubleshooting steps should you try next? (Choose three.)
    - **A.** Replace the battery in the stylus.
    - **B.** Replace the tip of the stylus.
    - **C.** Replace the stylus sensor in the tablet.
    - **D.** Shut down and restart the tablet.

14. When printing to a laser printer, all you get is random characters on the page. What is the first thing to try?
    - **A.** Replace the formatter board in the laser printer.
    - **B.** Reseat the printer cable.
    - **C.** Reinstall the printer driver.
    - **D.** Perform calibration to align the printheads.

15. You are unable to charge an Android smartphone. You tried a second charger that is known to work, and the phone will still not charge. What should you do next?
    - **A.** Replace the smartphone.
    - **B.** Use the secondary auxiliary port to charge the smartphone.
    - **C.** Use a DC charger instead of an AC charger.
    - **D.** Try a different outlet.

16. You are connecting a laptop to a conference room's A/V system for a presentation using an HDMI cable. Your screen image appears on the displays, but when you play a video, there is no audio. When you unplug from the A/V system, the audio plays on your laptop. What should you do next?
    - **A.** Use a DVI cable instead of an HDMI cable.
    - **B.** Reboot the laptop and reconnect to the A/V system.
    - **C.** Check the audio output settings on the laptop.
    - **D.** Restart the A/V system and reconnect to the laptop.

17. When an inkjet printer is printing jobs, it's making a squealing sound. What is most likely causing this problem?
    - **A.** A failing stepper motor.
    - **B.** Improperly refilled ink cartridges.
    - **C.** The paper caliper is too thick.
    - **D.** A failing formatter board.

18. An office has a flat-screen monitor in the lobby that displays information for visitors. Recently, an employee changed the material in the slide that's always shown. Even though the new slide can be seen, the old slide material still appears faintly on the screen. What is the solution to the problem?

A. Shut down and restart the monitor.

B. Shut down the presentation software. Reopen it, delete the old slide completely, and then restart presentation mode.

C. Degauss the monitor.

D. Replace the monitor.

19. An inkjet printer consistently feeds two or three pieces of paper for every page it's supposed to print. Which of the following could be the cause of the problem? (Choose two.)

A. The paper moisture content is too high.

B. The tension in the pickup tray needs to be adjusted.

C. The printer is using the incorrect tray setting.

D. The printer is using the incorrect page orientation.

E. The printer is experiencing finishing issues.

20. The video projector mounted to the ceiling in a conference room periodically shuts down during the middle of presentations without warning. When it shuts down, it won't power back on for about an hour. What should you check for?

A. Ensure that the proper mounting bracket was used for the projector.

B. Look for dust buildup in the projector's air filter.

C. Ensure that the input source is set correctly.

D. Look at the display settings to change the auto power-off time setting.

Performance-Based Question

You will encounter performance-based questions on the A+ exam. The questions on the exam require you to perform a specific task, and you will be graded on whether or not you were able to complete the task. The following requires you to think creatively in order to measure how well you understand this chapter's topics. You may or may not see similar questions on the actual A+ exams. To see how your answers compare to the authors', refer to Appendix B.

Your network users are sending print jobs to the printer, but they are stacking up in the queue and not printing. The printer appears to be online and has paper. How would you stop and restart the print spooler in Windows 11?

Appendix

A

Answers to Review Questions

Chapter 1: Motherboards, Processors, and Memory

1. A. The spine of the computer is the system board, otherwise known as the motherboard. All other computer components plug into the motherboard. On the motherboard, you will find the CPU, underlying circuitry, expansion slots, video components, RAM slots, and various other chips.

2. C. DDR4 SDRAM is manufactured on a 288-pin DIMM, as is DDR5 (although the two are keyed differently and are not compatible). DIMMs with 224 pins and 296 pins do not exist. DIMMs with 240 pins are used for DDR2 and DDR3 SDRAM.

3. B. If a system requires ECC RAM, then only ECC-RAM will work. Non-ECC RAM will not have error-correcting capabilities. There is no parity bit on RAM (at least not anymore).

4. D. The ITX motherboard family consists of smaller boards that fit in standard or miniature cases and use less power than their larger counterparts. ATX are "standard" sized motherboards in today's world. AT is a legacy form factor that was bigger than ATX. There are no DTX motherboards.

5. D. The Core i5-13600K is a Raptor Lake-series processor, which uses the LGA 1700 socket. LGA 1150 sockets are for older i7s, and the LGA 1366 socket is for even older ones than that. The LGA 1200 socket is for 10th generation Intel Core processors.

6. B. Thermal paste and thermal pads are used to attach heat sinks and fans to processors. A heat sink usually gets attached to the processor using thermal paste, and then the fan is attached to the heat sink. Superglue would make it adhere, but it does not transfer heat properly—it would likely melt and cause a mess and processor failure.

7. B. Overclocking is when someone sets the speed of the processor to run faster than it was rated for. While overclocking can make a computer faster, it can also cause the processor to produce more heat. Hyperthreading allows a processor core to handle multiple processes simultaneously. Virtualization is when you create multiple virtual machines on a computer. Multicore means that the processor has more than one core—it's not a configurable option.

8. A. Liquid cooling systems require a reservoir of liquid and a pump to circulate the liquid. It's possible that the system could malfunction, causing a mess inside the computer. Thermal paste is used to adhere a heat sink to a processor. Heat sinks and fans do not use pumps.

9. C. ARM processors usually are smaller and generate less heat than do x64/x86 CISC processors; therefore, they are better options for mobile devices. LGA is a type of CPU socket.

10. B. Soft power is the feature whereby the front power button acts as a relay to initiate various system power changes, depending on the duration that the button is held. Programmable power, relay power, and hot power are not valid power options.

11. B, F. DIMMs used in desktop motherboard applications have one of four possible pin counts. SDR SDRAM is implemented on 168-pin modules. DDR SDRAM is implemented on 184-pin modules. DDR2 and DDR3 are implemented on 240-pin modules with different keying, and DDR4 and DDR5 DIMMs have 288 pins but different keying. Older dual-channel DIMM modules have 232 pins. Modules with 200 and 204 pins are used in the SODIMM line, and there are no modules with 180 pins.

12. C, D. For virtualization to work, both the CPU and the BIOS need to support it. The operating system does as well. Virtualization support may need to be manually enabled in the BIOS. RAM does not need to support virtualization, and neither does the motherboard (except for the BIOS).

13. D. Most motherboards have a jumper or similar momentary closure mechanism that will allow you to clear the CMOS memory of any user settings and cause the BIOS to use factory defaults, including no user or supervisor passwords.

14. A. The easiest solution that works to cool your CPU is to connect the 4-pin connector into the 3-pin header. The missing pin allows you to control the speed of the fan. Without it, the fan will run at top speed, which is fine, albeit a little noisier. There is no 4-pin to 3-pin power adapter. The heat sink alone should not be relied on for proper cooling of modern CPUs, and an extra chassis fan isn't designed to specifically help cool the CPU.

15. C. The PCIe 1.1 specification provided 250 MBps of throughput per lane per direction. With the 2.x versions of PCIe, this rate was doubled to 500 MBps. PCIe 4.0 doubles that twice more. As a result, each v2.0 lane is capable of a combined 4 GBps. An x16 slot consists of 16 lanes, for a total bidirectional throughput of 64 GBps.

16. C. The reset button causes the computer to return to nearly the same point it is in when you power it on, but without the need for power cycling. Using Restart in the Start menu does not reboot as deeply as the reset button. Hibernation is a power state that completely removes power after saving the contents of RAM to the hard drive; pressing the power button is required to resume the session in the same manner as starting the computer after a complete shutdown. The power button cannot be used as a method of restarting the system.

17. B. None of the options are required, but a UPS is by far the most helpful among the answers in that loss of power during this procedure can range from annoying to devastating.

18. A. The only viable option listed is to install two new 32 GB DDR4 RAM modules. Check the motherboard documentation first to make sure the system will support that much memory, though. DDR4 does not come in modules larger than 64 GB. DDR5 is not compatible with DDR4 motherboards.

19. C. Laptop motherboards are usually custom-built for the specific case. A replacement board should be obtained that is the exact same size as the one that failed.

20. B. PCIe is the newest and fastest slot for video cards. AGP and PCI are legacy technologies. SATA is for hard drives.

Chapter 2: Expansion Cards, Storage Devices, and Power Supplies

1. D. Although technically PCI could be used for graphics adapters, PCIe supports high-speed, 3D graphic video cards. PCIe offers better performance than older graphics adapters. USB can stream video, but it is not used for attachment of internal graphics adapters. NVMe is an SSD communications interface.

2. E. If the slot is M-keyed, that means it has a block in the M position. For a drive to fit into that slot, it must have a slot in the M position (called being M-keyed). A drive that is B-keyed only will not fit because its slot will not line up with the block in the connector. A B+M-keyed drive has slots in both the B and M positions, so it will fit as well.

3. B. A network interface card (NIC) is used as a communications device. It enables the computer to talk to other computers on a network. It does not provide multimedia, input/output (such as a USB port does), or storage capabilities.

4. D. To protect against a power supply failure, use a redundant power supply, which is basically two PSUs in one unit. More wattage will not help if the power supply fails. There are no voltage output selectors on PSUs. There are voltage input switches on some, but that wouldn't help in this case. A modular power supply has removable and attachable cables to support a variety of peripherals.

5. B. The fastest SSD communications interface is NVMe. SATA 10,000 rpm and 15,000 rpm refer to conventional hard drive spin rates, and conventional hard drives will always be slower than SSDs.

6. B. Blu-ray discs have a single-sided, single-layer capacity of 25 GB. A double-sided, single layer DVD can store about 9.4 GB, and a double-sided, double-layer DVD can store about 17.1 GB.

7. A. Hot-swappable devices can be removed while the power to the system is still on. Warm-swappable devices need to be stopped in the operating system before being removed. The term has nothing to do with the heat level of the device.

8. C. Power supplies and AC adapters use standard wall outlets for an input of AC voltage, which they convert to the DC voltages required by the components to which they supply power.

9. A, C, E, F. A PC's power supply produces +3.3VDC, +5VDC, +12VDC, and –12VDC from a 110VAC input.

10. D. PC power supplies accept alternating current (AC) as input and produce direct current (DC) for the internal components. Europe requires the voltage selector switch be set at the higher setting, if the power supply does not auto-detect voltage input. SATA drives most often use a specific power connector that is not compatible with the Molex connector used by PATA drives. ATX motherboards can have a 20- or 24-pin single power connector but an adapter is not always required.

11. C. Today's hard drives, regardless of their rpm, have standard internal power connections. The other options are valid concerns when installing an internal drive.

12. A. Memory cards such as SD cards were designed for devices such as digital cameras. Optical discs are larger and used mostly with laptops and larger systems. Flash drives are USB. Some cameras may have a USB port, but that is external storage and not internal storage. mSATA is a laptop expansion card technology.

13. A. NVMe is a communications standard and there are no NVMe slots. This drive needs to be installed into an open PCIe slot. PCIe hard drives do not require an additional data cable or power connector.

14. E. For a video card, the best standard currently available is PCIe, and the widest channel width is x16. There are no PCI x16 or PCI x128 slots, nor is there a PCIe x128 slot. AGP is an older video expansion standard that is now obsolete.

15. B. Personal computers do not have permanently installed power supplies. Like other electrical and electronic components, power supplies can and do fail on a regular basis. Permanently mounting a power supply to a chassis would be a disservice to the consumer. You should consider the cumulative power needs of your installed components, and you might have to obtain adapters and splitters if you do not have enough or you have the wrong types of connectors coming from the power supply.

16. C. The smallest of the options is microSD, at 15 mm × 11 mm. Perhaps "micro" gives it away, but that also makes it easy to remember. SD cards are 32 mm × 24 mm. CompactFlash cards are the largest, at 36 mm × 43 mm. MiniSD cards measure at 21.5 mm × 20 mm.

17. C. The fastest drive will be an NVMe M.2 drive. M.2, PCIe, and SATA are different connector types, so there are no M.2 PCIe drives or M.2 SATA drives. NVMe over SATA is limited to 600 MBps transfer speeds, whereas M.2 is not.

18. D. RAID 5 arrays require a minimum of three hard drives to implement. If one drive fails, the data is fine and the system can still function, albeit slowly. If more than one drive fails, the system will be inoperable. RAID 5 provides both fault tolerance and a performance improvement.

19. C, D. Most likely, the new card will be detected and function properly once the computer is booted up. However, the driver will need to be installed. Odds are the new card will be a different model from the integrated card, which will require a new driver. On some systems, you will need to go into the BIOS and disable the integrated adapter.

20. A. Power supplies are rated in watts. When you purchase a power supply, you should make sure the devices inside the computer do not require more wattage than the chosen power supply can offer. The voltage is fairly standard among power supplies, and it has nothing to do with the devices connected to the power supply. Amperage and resistance are not selling points for power supplies. Another thing that could vary is the number of power connectors—make sure the new power supply has enough to support all installed peripherals.

Chapter 3: Peripherals, Cables, and Connectors

1. B, E. Analog monitors require an analog signal. VGA is an analog video standard and will work. DVI-A is also analog. (DVI-I would work as well.) HDMI, DVI-D, and DisplayPort are all digital standards.

2. A. USB 2.0 has a maximum speed of 480 Mbps. USB 3.0 is 5 Gbps, USB 3.1 is 10 Gbps, USB 3.2 is 20 Gbps, and USB4 is up to 40 Gbps.

3. C. A yellow USB port is an "always on" charging port. It could be either USB 2 or USB 3. USB-C connectors are oval in shape vs. the rectangular older USB connectors.

4. D. As you lower the resolution, the maximum refresh rate allowed tends to increase. The refresh rate is most often expressed in cycles per second (Hz), not millions of cycles per second (MHz). You must usually select the refresh rate that you want from the display settings dialog boxes, not through the monitor's built-in menu system, although the monitor can often tell you the refresh rate that you're using. Finally, both the monitor and adapter must agree on the refresh rate you select. If either device does not support a particular refresh rate, such a rate cannot be used.

5. C. DB-9 connectors are most often used with serial cables and are not associated with hard drives. USB-C, eSATA, and SATA are examples of connectors used with hard drives.

6. A. For smaller electronic devices such as cameras, micro-USB and mini-USB connectors are common. Thunderbolt is primarily used for video. Molex is a power connector. Lightning is an Apple proprietary connector that has been replaced by USB-C on newer devices.

7. B, D. HDMI cables transmit digital video and audio signals. They can be used to connect to projectors or monitors. HDMI cables can have standard HDMI connectors, mini-HDMI connectors, or micro-HDMI connectors.

8. A. Intel and Apple collaborated on Thunderbolt to add PCIe to VESA's DisplayPort and to make the resulting interface smaller and less expensive to connect. HDMI, DVI, and VGA are not equivalent to DisplayPort with PCIe added in.

9. B, C. USB4 is based on the Thunderbolt 3 specifications, and both natively use USB-C connectors.

10. C. Such a connection should not be made. DVI-I cables act like universal cables; they can connect two DVI-A interfaces or two DVI-D interfaces with adapters. Natively, they are used to connect two DVI-I interfaces, both of which are configured as either analog or digital. They are unable to convert the analog signal to a digital one, however. Analog and digital DVI interfaces are too disparate to interconnect.

11. D. A Molex connector is a power connector for IDE (PATA) storage devices. This includes hard drives and optical drives.

12. D. Thunderbolt 3 uses the oval-shaped USB-C connector. Thunderbolt 2 used a proprietary connector that was more of a square. Micro-USB and mini-USB connectors look more like trapezoids.

13. A, C, E. USB cables and eSATA cables can transmit power and data. You will find Lightning, USB-C, and micro-USB connectors on the ends of USB cables. A Molex connector is power-only. A DB-9 is for serial cables and is data-only.

14. D. The eSATA standards limit throughput to 6 Gbps, which is much slower than USB4 and Thunderbolt 3, both of which can get to 40 Gbps. Standard eSATA cables do not provide power, although there are versions of eSATA that do include power. The connectors are larger than USB-C connectors.

15. B. Peripherals most often use USB today, but they can also use D-sub connectors (rarely), and audio devices such as microphones and speakers use 1/8″ jacks. Molex connectors are power connectors for PATA hard drives and optical drives.

16. C. A projector uses a standard video cable. In this case, HDMI is the most likely choice. eSATA is for hard drives, and DB-9 is a slow serial cable. It's possible that the projector has a USB interface, but the best selection is HDMI.

17. B. Starting with the iPhone 5 and ending with iPhone 15, Apple iPhones used the proprietary Lightning connector for charging and connecting peripherals. Thunderbolt connectors are used on MacBooks for peripherals. iPhones do not have built-in mini- or micro-USB connectors.

18. B. Classic serial ports and serial cables have the slowest data transmission speeds of any cable discussed in this chapter.

19. C. Full Speed USB is 12 Mbps. High Speed is 480 Mbps. SuperSpeed runs at 5 Gbps and SuperSpeed+ at 10 Gbps.

20. B. VGA signals are analog, uncompressed, component signals that carry all of the video information for all three components of the original RGB signal. VGA signals are not digital, composite, or compressed.

Chapter 4: Laptop and Mobile Device Hardware

1. D. Laptop service manuals can be obtained from the manufacturer's website. It's very rare that paper service manuals are shipped with the laptop. Pressing F1 while in Windows will open Windows Help, and pressing F2 on many laptops during the system boot will take you into the BIOS/UEFI.

2. B. The inverter, which transforms DC power into AC power for an LCD backlight, is found in the display. It's not underneath the keyboard, next to the heat sink and fan, or near the processor, which are all located in the bottom half of the laptop.

3. A, B. The components of an LCD screen are the inverter, screen, and backlight. LCD display units will also have a Wi-Fi antenna (not the card), webcam, and microphone. There is no backdrop component.

4. A. The trackpad sits below the keyboard, where a user's palms may rest when he is typing. That contact may cause the mouse cursor to appear to jump on the screen. Disabling the touchpad should prevent this from happening.

5. B. Tethering is connecting a device to an Internet hotspot. One assumes in this case that the smartphone can function as a hotspot, and then the laptop will tether to it either via Wi-Fi or a wired connection. Internet connections cannot be shared over NFC or Bluetooth.

6. C. DDR4 SODIMMs have 260 pints. DDR2 SODIMMs have 200 pins, DDR3 SODIMMs have 204 pins, and DDR5 has 262. No SODIMMs currently have 288 pins, but DDR4 DIMMs do.

7. C. If the touchscreen is not responding properly, the first step is to recalibrate it. If it continues to have problems, the screen may be defective and might need to be replaced. Degaussing is a process used on old CRT monitors and is not relevant to touchscreens.

8. C, D. Two wireless connectivity options are NFC and Bluetooth; hotspot (not mentioned) is the third. USB and Lightning are wired connectivity methods.

9. C. Sometimes the only way to ensure that equipment doesn't walk away is to lock it down physically. Laptops come equipped with holes for cable locks that can be used to secure them to a desk or other workstation. GPS tracking might help after the device is stolen, but it's best to prevent that in the first place. A docking station does not prevent physical theft, nor does a screensaver with a password.

10. C. The laptop power supply converts AC power into DC power for the laptop's internal circuitry. The display backlight, if used, needs AC power. The inverter's job is to convert (invert) the DC power into AC power for the backlight.

11. D. To migrate data from one hard drive to another, you can manually copy the data or use a migration app. There is no NFC disk transfer utility.

12. C. The user needs a digitizer, which takes input from the device, such as drawn images, and converts them into electronic images on the computer. The user could also use a touchscreen for the same purpose. An inverter changes DC power into AC power for a display backlight. There is no capturer laptop device. A trackpad is a pointing device that allows a user to control the mouse cursor.

13. A. Laptops most often use SODIMMs (or possibly MicroDIMMs), which are much smaller than standard DIMMs used in desktops. DDR2 and DDR3 SODIMMs are 67.6 millimeters wide, whereas DDR4 SODIMMs are 69.6 millimeters wide. It's likely that the technician has DDR4 SODIMMs and the motherboard supports DDR3 SODIMMs.

14. A. A docking station made specifically for its associated brand and model of laptop can host desktop components permanently, regardless of whether the laptop is attached to the docking station. When the laptop's portability is not required, but instead use of the desktop components is the priority, attaching the laptop to the docking station makes such components available to the laptop without separately attaching each component.

15. D. The processor can reduce how fast it's working, which is called throttling, to help conserve battery life. CPUs do not perform underclocking, cooling, or disengaging.

16. B. Think of wattage as a "bucket" of power that the attached device can draw from. A bigger bucket simply holds more power but does not force the power on the device. Less wattage is not advised, however. Voltage can be thought of as the pressure behind the power to the device. Anything but the proper voltage is dangerous for the device. When you replace a laptop's AC adapter, you should match the voltage ratings of the original adapter. This also means that you should use an adapter with a fixed voltage if that matches the characteristics of the original; otherwise, obtain one that automatically switches voltages at the levels needed.

17. C. Battery calibration for Li-ion batteries allows the powered device to drain the battery's power before recharging. Battery exercising is the initial charging and discharging of nickel-based batteries so that they will function as expected. You should never short a battery's terminals, and replacement is a last resort, used when any battery has reached the end of its life.

18. D. Laptop hard drives commonly have a 2½″ form factor, and smaller drives, such as M.2, are becoming more popular. The most common form factor for desktop hard drives is 3½″. Laptop hard drives use the same drive technologies as their desktop counterparts, such as SATA. As with desktop hard drives, laptop hard drives are available in both solid-state and conventional varieties. Unlike desktop hard drives, laptop hard drives do not have separate power connectors.

19. B. A USB Bluetooth adapter should do the trick. There are Mini PCIe Bluetooth adapters, but installing one is trickier than inserting a USB device. There aren't really any Wi-Fi compatible headsets on the market—headsets are exactly what Bluetooth was designed for. Replacing the laptop is a much more expensive proposition.

20. C. A docking station is a device that a laptop can be plugged into that replicates common laptop ports, and may also have full-sized bays for PC desktop expansion cards and storage devices, such as hard drives and optical drives. A port replicator will replicate laptop ports but does not usually have expansion bays. There is no specific laptop base accessory or a Mini PCIe dock. (Mini PCIe is an internal laptop technology.)

Chapter 5: Printers and Multifunction Devices

1. A. If a user shares their printer on a network, their computer becomes a print server. The shared printer itself is the printer share. It's possible other users could use the device as a network scan service but that's not specified. An automatic document feeder (ADF) is a physical feature of some MFDs.

2. B, C. A page printer is a type of computer printer that prints a page at a time. Common types of page printers are the inkjet printer and the laser printer.

3. D. The rate of transfer and the ability to recognize new devices automatically are two of the major advantages that currently make USB the most popular type of printer interface. However, it is the network printer interface that allows the printer to communicate with networks, servers, and workstations.

4. D. Dot-matrix printers are impact printers and therefore are the best for multipart forms. Daisy-wheel printers can be used with multipart forms as well.

5. A. The exposing step uses a laser to discharge selected areas of the photosensitive drum, thus forming an image on the drum.

6. C. An automatic document feeder (ADF) allows users to put multiple pieces of paper into it, and scan them in all as one job. SMB is used with network scan services to scan to a folder. An MFD is another name for a printer/scanner combination device. Flatbed scanners only allow for one piece of paper at a time.

7. A. Audit logs will track print jobs and the users who printed them. User authentication is where a user is required to log into the printer. This can be done via badging—or swiping an employee badge on a badge reader on the printer—or by manually entering credentials. Secured prints is when a print job is held until a user authenticates at the printer and tells it to start printing.

8. B. In an inkjet printer, the ink cartridge is the actual printhead. This is where the ink is expelled to form letters or graphics. Toner cartridges are used by laser printers to store toner. A daisy wheel is the device that impacts the letters on the paper in a daisy-wheel printer. Paper trays are the storage bins in laser printers and inkjet printers that allow the pickup rollers to feed the paper into the printer.

9. D. The correct sequence in the EP imaging process is processing, charging, exposing, developing, transferring, fusing, and cleaning.

10. C, E. Two ways for users to authenticate are to use user authentication (which requires typing in the username and password) or badging, which requires the MFD to have a badge scanning feature. Secured print will save print jobs until they are ready to be released. SMB is a file and print sharing protocol used by MFDs to send scanned documents to a shared network folder. There are no printer authentication logs, but audit logs may be used to see who has used the MFD.

11. A, B, D. An electrophotographic (EP) laser printer toner cartridge includes the toner, print drum, and cleaning blade. The laser is usually contained within the printer, not within the toner cartridge.

12. A. After a laser has created an image of the page, the developing roller uses a magnet and electrostatic charges to attract toner to itself and then transfers the toner to the areas on the drum that have been exposed to the laser. The toner is melted during the fusing stage. The laser creates an image of the page on the drum in the exposing stage. An electrostatic charge is applied to the paper to attract toner in the transferring stage, which happens immediately after the developing stage.

13. A, C, D. Printers can communicate via parallel, serial, USB, wireless, and network (Ethernet) connections. SATA is a bus interface that connects host bus adapters to storage devices such as hard drives and optical drives.

14. C. If a printer is using out-of-date or incorrect printer drivers, the printer may produce pages of garbled text. The solution is to ensure that the most recent printer drivers are downloaded from the manufacturer's website.

15. B. The daisy-wheel printer gets its name because it contains a wheel with raised letters and symbols on each "petal."

16. A. The high-voltage power supply (HVPS) is the part of the laser printer that supplies the voltages for charging and transferring corona assemblies.

17. C. The transfer corona assembly gets the toner from the photosensitive drum onto the paper. For some printers, this is a transfer corona wire; for other printers, it is a transfer corona roller.

18. D. Developing happens after exposing. The correct order is processing, charging, exposing, developing, transferring, fusing, and cleaning.

19. B. The fuser assembly presses and melts the toner into the paper. The transfer corona transfers the toner from the drum to the paper. The printer controller circuitry converts signals from the PC into signals for the various printer assemblies. The paper transport assembly controls the movement of the paper through the printer.

20. A. Firmware upgrades for laser printers are downloaded for free from the manufacturer's website. A technician does not need to install a new chip, because firmware is upgraded via software. It's unlikely that the manufacturer will send you the upgrade on an SSD; it will refer you to its website to download it.

Chapter 6: Networking Fundamentals

1. D. Companies that want to ensure the safety and integrity of their data should use fiber-optic cable because it is not affected by electromagnetic or radio-frequency interference. Even though some copper cables have shielding, they are not immune to EMI or RFI. This eliminates twisted pair and coaxial. CSMA/CD is an access method, not a cable type.

2. C. Coaxial cable can use BNC connectors and F-type connectors. RJ11 and RJ45 are twisted pair connectors. SC is a fiber-optic connector.

3. C. A crossover cable will flip pins 1 and 3 on each end. Therefore, this end needs to have the white/orange wire as pin 1. The correct order for this cable should be white/orange, orange, white/green, blue, white/blue, green, white/brown, brown.

4. B. Carrier Sense Multiple Access with Collision Detection (CSMA/CD) specifies that the NIC pause before transmitting a packet to ensure that the line is not being used. If no activity is detected, then it transmits the packet. If activity is detected, it waits until it is clear. In the case of two NICs transmitting at the same time (a collision), both NICs pause to detect and then retransmit the data. CSMA/CA listens to the line as well but tries to avoid collisions. T568A/T568B are twisted pair wiring standards. Demand priority is a media access method, but we didn't discuss it at all in the chapter as it's not related to the A+ exam. It's used with a standard called 100BaseVG.

5. A. Cat 5 was the first UTP standard ratified for 100 Mbps transmissions. Before that, Cat 3 was popular but only supported 10 Mbps. Cat 5e, Cat 6, and Cat 6a all support at least 1Gbps and came after Cat 5.

6. D. Direct burial cable is designed to be buried underground, as its name implies. Direct burial can be fiber-optic or copper. Plenum cable is used for cables run through spaces that circulate breathable air inside a building.

7. C. Fiber-optic cable can span distances of several kilometers because it has much lower attenuation, crosstalk, and interference as compared to copper cables. Single-mode fiber (SMF) can run longer distances than multimode fiber (MMF).

8. A. Routers are designed to route (transfer) packets across networks. They are able to do this routing, determining the best path to take, based on the internal routing tables that they maintain.

9. D. A managed switch can create virtual LANs (VLANs) using its management software. A firewall is a hardware or software solution that serves as your network's security guard. Patch panels and unmanaged switches cannot create VLANs.

10. A. Bluetooth networks are often called wireless personal area networks (WPANs).

11. C. Fiber-optic Internet connections use an optical network terminal (ONT) to terminate the fiber and translate it into a digital connection over copper cable inside your home. A cable modem and DSL modem connect copper to copper. Power over Ethernet (PoE) is not relevant in this situation.

12. A. For areas where a cable must be fire retardant, such as in a drop ceiling, you must run plenum-grade cable. Plenum refers to the coating on the sleeve of the cable, not the copper or fiber within the cable itself. PVC is the other type of coating typically found on network cables, but it produces poisonous gas when burned. Direct burial cables are designed to be buried underground.

13. D. The local connector (LC) is a mini form factor (MFF) fiber-optic connector developed by Lucent Technologies. If it helps, think of LC as "Little Connector."

14. A. If you need to make a connection that is 5 kilometers long, then you are limited to fiber-optic cable, specifically single-mode fiber (SMF). Two common SMF standards are 10GBaseER and 10GBaseEW. (Think of the *E* as *extended*.) A *T* designation is an Ethernet standard that refers to twisted pair. SR (short reach) and DR (distance reach) are fiber standards that do not stretch for 5 kilometers.

15. B, C. Storage area networks (SANs) have many advantages, including block-level data storage, which is more efficient. They also reduce network loads, provide faster data access, and area easily expandable. They do not inherently provide more storage space (although they are easy to expand) or do automated data backups (although they are easier to set up).

16. C. A straight tip (ST) connector uses a twist to lock mechanism. BNC twists to lock as well, but it's for copper cable. SC and LC snap to lock into place.

17. B. The job of a firewall is to block unwanted network traffic. Firewalls do this by using a list of rules called an access control list (ACL). Routers connect networks to each other, and switches are central connectivity points for networks. A network attached storage (NAS) device is like a dedicated file server.

18. A. The two RG standards used for cable television are RG-6 and RG-59. Of the two, RG6 is better because it can run longer distances and can handle digital signals. RG-59 is for analog signals only. RG-8 is thicknet coax, and RG-58 is thinnet coax.

19. D. A switch works at Layer 2; therefore, it works with hardware (MAC) addresses. Hubs and patch panels work at Layer 1 and deal with bits, and routers work at Layer 3 and work with logical IP addresses.

20. B. A punchdown block connects bare wires to its frame to make a network connection. There are no other cable connectors used on that end of the cable. A patch panel uses RJ45 or other connectors, as does a PoE injector. Access points are wireless and therefore don't use wires.

Chapter 7: Introduction toTCP/IP

1. B. IPSec is the only one of the protocols listed that works at OSI layer 3, and therefore the only one that can encrypt all communications. HTTPS uses TLS or SSL, which operate at OSI Layer 7.

2. B. Lightweight Directory Access Protocol (LDAP) is designed to access information stored in an information directory typically known as an LDAP directory or LDAP database. This typically includes employee data and network configuration data. FTP is designed to transfer files. RDP is a remote desktop protocol, and CIFS is used to share files and printers.

3. A. HTTP uses port 80. HTTPS uses port 443; Telnet uses port 23; and POP3 uses port 110.

4. D. An IPv6 interface is not limited in the number of addresses that can be assigned, although there could be limitations based on practicality.

5. A, B, C. An IPv6 address contains 128 bits, written in eight 16-bit fields represented by four hexadecimal digits. Option A contains all eight fields expressed in full. Option B is an IPv4 address expressed in IPv6 form. Option C is the same address as option A, but written in accepted shorthand. Option D is not valid because the double colons (::) can be used only once within an IPv6 address.

6. A. DNS servers resolve hostnames to IP addresses. On the Internet, a DNS server needs to have a public IP address. The address 10.25.11.33 is in a private address space, so that address would not be valid for a DNS server on the Internet.

7. D. The address assigned to the computer is an APIPA address. Microsoft client computers (and others) will configure themselves with an address in this range if they are unable to reach a DHCP server.

8. C. Simple Mail Transfer Protocol (SMTP) is responsible for sending email. IMAP4 and POP3 both receive email. SNMP is a network management protocol.

9. D. Remote Desktop Protocol (RDP) works on port 3389. DNS works on port 53; IMAP4 works on port 143; and LDAP works on port 389.

10. D. A VLAN is a virtual LAN configured on a switch, and breaks logical networks into smaller broadcast domains, reducing the number of broadcasts each computer will receive. DNS resolves hostnames to IP addresses. DHCP assigns IP configuration information. A VPN is a virtual private network over a public Internet connection.

11. A, D. TCP is a connection-oriented protocol that establishes virtual circuits and acknowledges delivery of packets. Because of these features, it has higher overhead than UDP and is a little slower.

12. B. The HTTP protocol is inherently unsecure, but the HTTPS protocol is secure. (SSH and TLS are secure as well, but they are not protocols used to connect to websites.)

13. A. The router is your doorway out into other networks, and it is known in TCP/IP terms as the default gateway. Without this configuration option, you would not be able to get to external networks. The subnet mask tells a host which network it's on. The DNS server resolves hostnames to IP addresses, and a DHCP server provides IP configuration information to hosts.

14. A, B. The only mandatory IPv4 configuration items are an IP address and a subnet mask. If you are not connecting to another network, you do not need a default gateway. DNS servers resolve hostnames to IP addresses, but they are not mandatory.

15. D. IMAP4 and POP3 are the two protocols that are used for email retrieval. Of the two, only IMAP4 provides security features. SMTP sends email. LDAP is a directory services protocol.

16. D. A virtual private network (VPN) is a secure point-to-point connection made over a public network. NAT translates private IP address to public IP addresses. APIPA is an automatic IP address assignment if a client computer can't reach a DHCP server. VLANs are virtual local area networks configured through switches.

17. A. Common Internet File System (port 445) is used to share files and printers on a network. Remote Desktop Protocol (port 3389) is used for remote computer administration.

18. E. IPv6 does not have broadcasts. IPv6 does have multicasts, which are a bit like targeted broadcasts. ::1 is the loopback address used for troubleshooting. FE80:: is the first part of a link local address. FF00:: is the first part of a multicast address. :FFFF will often be used as part of an IPv4 address designation (the IPv4 address will follow it).

19. A. Lightweight Directory Access Protocol (LDAP, port 389) is the "phonebook" protocol of TCP/IP, helping manage user information. Simple Mail Transfer Protocol (SMTP, 25) is an email-sending protocol. NetBIOS over TCP/IP (NetBT, 137-139) allows computers to communicate with each other over the network using APIs. Server Message Block (SMB, 445) is a file- and printer-sharing protocol.

20. C. The Address Resolution Protocol (ARP) resolves IP addresses to hardware (MAC) addresses. RARP does the reverse—it resolves MAC addresses to IP addresses. DNS resolves hostnames to IP addresses. DHCP automatically configures TCP/IP clients.

Chapter 8: Network Services

1. B. For secure transactions, the web server will be using HTTPS, which uses port 443. If the website works from an internal workstation, then the server is fine. It's most likely that the firewall is blocking inbound traffic on port 443. Port 23 is Telnet.

2. C. It sounds like the manager wants a unified threat management (UTM) device. They are designed to be one-stop network protection devices. Spam gateways will help with spam email but not other malware. A load balancer spreads work around to multiple servers. A proxy server fulfills requests for clients.

3. C. If the data on the server does not need to be accessed via the Internet, then the server should be in the most secure place possible, which is inside the firewall(s) in the secure network. The screened subnet is also behind a firewall, but by nature will be partially open to the outside world, so it's not as secure as the internal network.

4. A, C. Print servers should make printers available to clients and accept print jobs. They also process print jobs and manage print priorities. Finally, they provide client computers with the right print drivers when the clients attempt to install the printer. They do not turn printers off on demand nor provide notification that a job has printed.

5. A. A proxy server can be configured to block access to websites that contain potentially dangerous or inflammatory material. Web servers host web pages, some of which may have objectionable content. DHCP servers provide clients with IP addresses, and DNS servers resolve host names to IP addresses.

6. B. Every DHCP server needs to have a scope, which is the range of addresses available to clients, as well as other options that it can give to client computers. A DHCP server can optionally provide clients with the address of a DNS server. There is no DHCP range. A DHCP relay agent is a system configured on a subnet with no DHCP server that relays DHCP requests to the DHCP server.

7. A. DNS server records are contained in the zone file, which must be configured by administrators. A hosts file is an alternative to using DNS (but that does not work well when scaling to the Internet). A scope is created on DHCP servers. There is no DNS proxy.

8. D. Simple Mail Transfer Protocol (SMTP) is used to transfer (send) email between servers. POP3 and IMAP4 are used to download (receive) email. SNMP is Simple Network Management Protocol and not related to email.

9. D. A load balancer can spread the work around to multiple servers. It accepts the inbound request and then sends it to the most appropriate web server. DNS resolves host names to IP addresses. DHCP provides IP configuration information. A proxy makes requests (usually outbound) on behalf of clients.

10. A, B, C. Legacy systems are ones that use older hardware, software, or network protocols that are not commonly used today. A system with only 1 GB of RAM might be woefully underpowered, but that in and of itself does not make it a legacy system.

11. D. The best answer is to create a scope with the available addresses that the server will manage. Then, if you want the server to assign the same address each time to servers, printers, and router interfaces, create reservations for these addresses. An exclusion could work, but then you would have to manually assign addresses to those devices and the DHCP server wouldn't assign all addresses.

12. C. IPv6 hosts have AAAA records in a DNS zone file. IPv4 addresses have A records. A CNAME record is an alias, and TXT is a text entry.

13. B. The purpose of an NTP server is to provide the correct time to all computers (or devices) on a network. A local NTP server can query national severs but does not need to; it can function on its own. NTP has nothing to do with newsfeeds or SCADA.

14. A, B. Unified threat management (UTM) systems can include antispam capabilities. DMARC is a type of spam management inclusion into DNS zone files. SCADA refers to legacy hardware systems, and AAA is authentication, authorization, and accounting.

15. B. Authentication refers to verification of user identity. Authorization is the set of permissions a user has to access resources. Accounting is the tracking or logging of who accessed what resource.

16. C. A syslog server collects and tracks error messages sent from network devices such as routers. AAA is for authentication, authorization, and accounting. NTP is a time server. UTM is unified threat management, which is a security server or servers.

17. D. Supervisory control and data acquisition (SCADA) systems are legacy systems that often control public infrastructure. DKIM, SPF, and DMARC all deal with spam management.

18. A. If a company wants to implement DHCP, it only needs one DHCP server. In fact, having more than one can cause administrative headaches if they are not configured properly. If a company wants to run their own DNS, they are required to have two DNS servers.

19. C. Smart assistants are examples of Internet of Things (IoT) devices. UTM is unified threat management, which is a security service. SCADA refers to legacy hardware and software. AAA is authentication, authorization, and accounting.

20. A. Sender Policy Framework (SPF) is the simplest of the three DNS TXT record types that deals with spam management. It authenticates servers based on their IP address. DomainKeys Identified Mail (DKIM) authenticates servers using a public-private key pair. Domain-based Message Authentication, Reporting, and Conformance (DMARC) allows domains to decide what to do to emails that fail SPF or DKIM authentication. SCADA refers to legacy hardware and software.

Chapter 9: Wireless and SOHO Networking

1. B, D. Both 802.11b and 802.11g operate in the 2.4 GHz range and use similar transmission standards. Alternatively, 802.11a and 802.11ac operate in the 5GHz range.

2. C. A service-set identifier (SSID) is the unique name given to the wireless network. All hardware that is to participate on the network must be configured to use the same SSID. Essentially, it is the network name. When you are using Windows to connect to a wireless network, all available wireless networks will be listed by their SSID.

3. A, C. The two technologies that 802.11ac employs to achieve high throughput are channel bonding and MIMO. Channel bonding is the combination of multiple smaller channels into one large channel for greater bandwidth. MIMO is enhanced over 802.11n to allow for more multiple inputs and outputs. 802.11ac also uses beamforming, but that helps the range, not the throughput.

4. D. WEP was the original encryption standard developed for Wi-Fi networks, but it is easily hacked. WPA is an upgrade, but WPA2 is more secure and incorporates the entire 802.11i standard. WPA3 is of course the safest.

5. C. Wi-Fi 6 is faster than Wi-Fi 5, and the network range is similar. More users should be able to access each access point, not fewer. Devices may have lower power usage in Wi-Fi 6 due to a new sleep feature.

6. A, B. You should always change the default administrator name and password as well as the default SSID when installing a new wireless router. Enabling encryption is also a good idea, but WPA, WPA2, and WPA3 are better options than WEP. The channel has nothing to do with security.

7. D. Of the options listed, fiber provides the best speed. DSL and cable provide similar speeds, but can't match fiber. Satellite will be the slowest of the options listed.

8. D. Network Address Translation (NAT) allows users to have a private IP address and still access the Internet with a public IP address. NAT is installed on a router and translates the private IP address into a public address for the user to access the Internet. DHCP assigns IP configuration information to clients. DNS resolves hostnames to IP addresses. DSL is a type of broadband Internet access.

9. C. There are 14 communication channels in the 2.4 GHz range, but only the first 11 are configurable. The three nonoverlapping channels are 1, 6, and 11.

10. A, B. Two standards listed—802.11be (46 Gbps) 802.11ac (6.9 Gbps)—give users the required throughput. 802.11ax would do so as well, although it's not an option. The 802.11g (54Mbps), and 802.11n (600Mbps) standards are too slow.

11. B. To join the network, client computers need to find the SSID, ensure that the security settings are correct (to match the router), and enter the security passphrase. As an administrator, you should have configured this passphrase to be different from the router's administrator password.

12. B. Radio frequency identification (RFID) can operate in three frequency bands: 125 kHz to 134 kHz, 13.56 MHz, or 856 MHz to 960 MHz. Bluetooth is 2.4 GHz; NFC is 13.56 MHz; and LTE is between 600 MHz and 6 GHz.

13. B. The feature of Wi-Fi 6 that reduces channel interference is Basic Service Set (BSS) coloring, which adds a field to the wireless frame that distinguishes it from others. Orthogonal Frequency Division Multiple Access (OFDMA) is a modulation technique that improves speed because it allows transmission to multiple clients at once. MU-MIMO also increases speed because in Wi-Fi 6 it works for uplink connections as well as downlink ones. Dynamic Frequency Selection (DFS) was introduced with 802.11n, and it detects radar interference and adjusts frequencies accordingly.

14. B. WEP could use a 64-bit or 128-bit security key, but it was a static key. TKIP introduced a dynamic per-packet key. AES and CCMP came after TKIP.

15. B. Wi-Fi 6 (802.11ax) was the first standard to support WPA3 security and OFDMA modulation.

16. D. MAC filtering is a security option that can specify that only computers with specific MAC (hardware) addresses can access the network. Port forwarding is a feature of firewalls. NAT translates private IP addresses into public ones. SSID is the wireless network name.

17. B. The set of rules for access on a firewall is called an access control list (ACL). A service level agreement (SLA) is an agreement on service level for quality of service (QoS). NAT translates private IP addresses into public ones. APIPA assigns an IP address to a client that can't reach the DHCP server.

18. C. If your router is using AES, the clients will need to use WPA2 or WPA3. TKIP is a protocol utilized by WPA. WEP is the weakest of the encryption options.

19. A. Bluetooth also operates in the 2.4 GHz range. Satellite uses satellite waves; 5G cellular mmWave is between 24 GHz and 86 GHz; and NFC uses 13.56 MHz.

20. D. The gateway, also known as a default gateway, is the address to the router to take a client to the next network (or Internet). Gateways can be configured with a static assignment or dynamically through a DHCP server. APIPA will automatically configure a client with IP address information if it's set to get its configuration form a DHCP server but none is available. APIPA-configured clients can't get on the Internet.

Chapter 10: Mobile Connectivity and Application Support

1. D. Whenever you have a reliable Wi-Fi signal available, which you do not pay for "as you go," you should feel free to disable your cellular access to data networking. The wireless network is often faster and does not cost you anything to use. If data networking works better when you leave the cellular network available as a fallback, you can choose to disable it only for certain large downloads or disable it completely until the download is complete. Phone calls will continue to be sent over the cellular network. This does not imply that your data is not using an available Wi-Fi connection. Therefore, unless you are particularly sensitive to the situation or know that data is going out over your cell access, you might not need to disable the cellular data-networking feature.

2. C. The range of Bluetooth connections is considered short compared to Wi-Fi connections. Bluetooth is a fully standardized protocol that supports file transfers using FTP. Rebooting, or even restarting, paired devices is not a requirement for Bluetooth connection.

3. A. SMTP sends mail to the server on TCP port 25. POP3 uses port 110. IMAP4 uses port 143. POP3 uses port 995 securely over SSL or TLS.

4. B. SMTP with TLS uses port 587 by default. With SSL, SMTP uses port 465. IMAP4 with SSL/TLS uses port 993, whereas POP3with SSL/TLS uses port 995.

5. B. Four commercial email providers to be familiar with are iCloud, Google/Inbox, Exchange Online, and Yahoo Mail. Hotmail used to be a commercial email provider, but it no longer exists; it's now part of `Outlook.com`.

6. C. In iOS (at least iOS 18), location services are under Settings ➢ Privacy & Security. There are no location services, communications, and GPS sections directly in Settings.

7. C. Each mobile phone has a separate processor that manages wireless communications, also known as radio communications or baseband communications. There is a baseband OS that manages this, and it works with the primary OS on the phone. A SIM OS is for managing data transfers between the phone and the SIM chip. There is no wireless OS.

8. A. Mobile application management (MAM) is a service that allows administrators to control corporate software on mobile devices, including updating and wiping apps. Mobile device management (MDM) controls which devices are on the network and can wipe entire devices. Location services does not manage software. Synchronization can back up apps and data but does not help with updates or deleting apps

9. A. Microsoft 365 is able to sync desktop settings between multiple Windows-based computers. ActiveSync is for syncing mobile devices with an Exchange email server. Mobile application management (MAM) and mobile device management (MDM) are for managing apps and network access for mobile devices. iCloud is a cloud-based storage solution.

10. C. The next step is to enter the PIN code. The steps, in order, are as follows: enable Bluetooth, enable pairing, find a device for pairing, enter the appropriate PIN code, and test connectivity.

11. B. The subscriber identity module (SIM) is a chip in a mobile device that stores user information, phone identification, and security data. It does not store user data, manage wireless connections, or location services.

12. B. GPS is generally more accurate than cellular when determining location. eSIM is a SIM chip built into a mobile device's motherboard. MDM is a mobile device management concept.

13. B. Mail access uses standard secure or unsecure TCP ports, not UDP ports. Exchange access to such mail services is not unheard of, but it is exceedingly rare. Mobile devices tend to have email clients built-in.

14. A, B, C. On an Android phone, airplane mode disables all radio communications, including cellular, Wi-Fi, and Bluetooth (and NFC as well). It does not affect the opening of any other apps or the lock screen.

15. A, C, D. Three common connection types for synchronization to a desktop (or laptop) are cellular, Wi-Fi and USB. Bluetooth is not used for desktop synchronization. Tethering is when a laptop or other device is granted cellular network access by connecting to a mobile device, such as a smartphone.

16. C. The user has most likely reached their cellular data cap. Cellular connections (including data) work fine on commuter trains. If the cellular connection were disabled, the user would not be able to call using the device. BYOD versus corporate configurations will not affect the ability to download data.

17. A. The subscriber identity module (SIM) chip contains the information needed to join the cellular network of the provider. If a device has a SIM chip, it can be replaced. eSIM is built into the motherboard and electronically reprogrammed. GPS is a location service. BYOD is bring your own device, a type of device configuration for using a corporate network.

18. B. Mobile device management (MDM) is a type of service that allows network administrators to specify which mobile devices can join the network. It also gives administrators the ability to remotely lock and completely wipe mobile devices. Mobile application management (MAM) lets an administrator control the corporate apps on a mobile device. An eSIM has phone and user configuration information. Location services can help find a mobile device, but not remotely wipe it.

19. C, D. Apple devices natively sync with iCloud and iTunes. Android devices sync with Google Drive. SharePoint is a Microsoft service which can be used to share files.

20. B. GPS, cellular, Bluetooth, and Wi-Fi can all be used as part of a location services solution. GPS is free, and Wi-Fi and Bluetooth likely are as well, if part of a company or organization's network. Cellular requires a contract with a provider, which is a paid service.

Chapter 11: Virtualization and Cloud Computing

1. A. Platform as a service (PaaS) is probably the right level of service for the developer team. It provides infrastructure, like IaaS, and also supplies needed programming elements. Infrastructure as a service doesn't provide runtime environments; think of it as hardware. Software as a service provides apps and is too high of a service for what's needed here. There is no current DaaS on the exam objectives.

2. A. The ability to expand services quickly refers to elasticity (or specifically, rapid elasticity). Availability guarantees uptime for services. All cloud services use resource pooling. Metered utilization is how many suppliers track usage and charge accordingly.

3. D. A hybrid cloud provides the best of public and private clouds. You get the scalability and cost effectiveness of a public cloud but also the security that you need for important files on the private portion of the cloud.

4. C. When multiple organizations with similar objectives want to combine efforts in a cloud, the best choice is generally a community cloud. This allows for the flexibility and scalability normally found in a public cloud, but it also limits the number of users to a smaller, trusted group. Hybrid clouds also combine public and private, but are for a single company to use.

5. B, C, D. Cloud solutions are great for enhancing scalability and reliability while generally lowering costs. Security could be an issue with cloud computing, depending on your organization's needs, because the resources aren't locally controlled and managed.

6. B. Each virtual machine will use its own virtual NIC, so you need three virtual NICs. The virtual NICs will communicate with one virtual switch managed by the hypervisor. The virtual switch will communicate with one physical NIC on the host system.

7. B. There needs to be enough RAM to support both OSs, so the answer is 6 GB. More is better, though!

8. A, B, C. Legacy systems are ones that use older hardware, software, or network protocols that are not commonly used today. A system with only 1 GB of RAM might be woefully underpowered, but that in and of itself does not make it a legacy system.

9. D. Each instance of the OS you are running requires its own security software.

10. B. A sandbox is a temporary operating system environment, kind of like a lite version of a virtual machine. It's ideal for testing software such as this. An AAA server is for authentication, authorization, and accounting of security. Supervisory Control and Data Acquisition (SCADA) is a hardware and/or software environment used to control industrial systems, and in many cases it is older, so therefore it is considered legacy. Application virtualization is used for legacy apps or cross-platform virtualization.

11. A. A Type 2 hypervisor sits on top of an existing OS, meaning that OSs installed in VMs will compete for resources with the host OS. The amount of resources available to a guest OS can be configured. Virtual OSs can get on the physical network if configured properly.

12. C. Containers are used to quickly deploy single apps and all of the files needed to run them, regardless of what platform (OS) the user is using. Hypervisors virtualize entire machines. VDI is the virtual desktop infrastructure.

13. B. Public clouds are usually the most cost effective. The client doesn't need to purchase any hardware, the cloud provider benefits from economies of scale.

14. B, D. Containers are faster to deploy than hypervisors. They can hold more than one app, but best practices usually specify that they hold only one. Containers do not have their own RAM requirements. They are designed to work on any OS.

15. D. Multitenancy happens all the time in cloud computing, meaning that multiple clients will share the same hardware resources. Their apps and data are separated though to avoid any data security issues.

16. B. A type 2 hypervisor is the best choice. It can sit on top of the developer's Windows workstation, and virtualize other OSs as needed for testing. A Type 1 (bare-metal) hypervisor is better suited for server virtualization. Containers may be useful in quickly deploying the apps once they are complete and ensuring they can run in multiple operating environments. VDI is the virtual desktop infrastructure.

17. A. Infrastructure as a service (IaaS) provides hardware, which includes storage space. Platform as a service (PaaS) would include software development tools, and Software as a service (SaaS) would include apps as well. Neither of those seem needed in this instance. VDI is virtual desktop infrastructure.

18. D. The virtual desktop infrastructure (VDI) is the hardware and software needed to create a virtual environment. Hypervisors are the software packages that create and manage virtual environments. Containers are used for software virtualization.

19. A. Ingress (uploading data) and egress (downloading data) may be part of metered utilization in a cloud agreement. Generally, ingress is free, but many cloud providers will charge for data egress. Egress is not related to containers or VDI.

20. C. If the CSP has guaranteed three clients 10 TB each, then the CSP needs at least 30 TB of available storage space. Clients may share hardware in a multitenancy environment, but each is guaranteed its contract terms. Having extra storage available for elasticity purposes is a good idea but not required.

Chapter 12: Core Hardware and Network Troubleshooting

1. D. This system is producing a beep code and not completing the POST process. Different BIOS/UEFI manufacturers use different beep codes. Look up this beep code on the manufacturer's website before replacing any hardware.

2. B. Capacitors store energy and are located on the motherboard and inside the power supply. Since the capacitors are exposed, it likely means the technician is seeing the problem on the motherboard. CPUs and RAM do not have capacitors.

3. B, D. Intermittent shutdowns are hard to troubleshoot but are most likely hardware-related. When troubleshooting, always make one change at a time to try to resolve the issue. Replacing the RAM or the CPU to see if it resolves the problem could be a good next step, but replacing three components at once is too many changes unless you have specific reason to believe it is all three. Reinstalling Windows is probably overkill as well. Narrow down the issue before implementing sweeping solutions such as that.

4. C. When troubleshooting, always question the obvious. In this case, the system beeps once, which is usually a signal from the BIOS that POST completed successfully, and air is coming from the system fan. All of that appears to be normal. Check to see that the monitor is connected properly to the desktop computer and that it's receiving power.

5. A, B. When the system is powered off, BIOS settings on a motherboard are maintained by the CMOS battery. If the battery fails, the system will likely forget the BIOS time and date settings and the boot order configuration. Replacing the CMOS battery will resolve the issue.

6. C, D. Continuous reboots are most likely caused by the motherboard or CPU failing. It's possible it's the RAM, but less likely. The power supply is the least likely cause of the problem in this scenario.

7. A, C. If the wall outlet is functioning and the power cable is connected properly (always check the obvious), then the most likely culprits are the power supply or the power cord. The motherboard, CPU, and RAM are unlikely to cause a complete power failure in a computer.

8. C. The processor runs the hottest of all computer components and is the most likely to overheat. Motherboards and RAM don't generate a lot of heat. Power supplies most often have fans to help dissipate heat, and don't run as hot as CPUs.

9. A. Only components with moving parts can make ticking or clicking noises. (Bombs can too, but hopefully you recognized that as a fake answer!) This question assumes that the hard drive is an HDD and not an SSD, but none of the other components would make a ticking noise.

10. A, B, C. Only components with moving parts can make ticking, clicking, or squealing noises. It's probably a fan of some sort (such as the CPU fan or power supply fan), or it could be a conventional hard drive. SSDs, CPUs, and RAM do not have moving parts and therefore could not be causing this problem.

11. A. Smoke and foul odors such as burning smells are not normal for computers! The most likely component causing the problem is the power supply. With the system off, open the computer case (not the power supply case!) to examine the inside for possible physical damage.

12. C. When Windows-based systems crash, as can happen due to a memory, CPU, or motherboard problem, they will most likely display a Blue Screen of Death (BSOD). The good news is that the BSOD can often provide useful information, such as an error code, to help troubleshoot the problem. There is no stop sign proprietary crash screen, and the pinwheel is used in macOS. A black screen may indicate a system crash but is not proprietary to one operating system.

13. C. Distance matters in wireless networks. Being too far away from an access point will result in a weaker signal, which can cause intermittent signal drops and poor performance. Moving closer to the wireless access point (WAP) or wireless router can help, as can installing additional WAPs or wireless routers. The IP address from the DHCP server won't matter, and since they reported no error message, there is not likely to be an IP address conflict. Unless there are other issues, replacing the NIC seems to be overkill.

14. B. The `ipconfig` command is perhaps the most-used utility in troubleshooting and network configuration. The `ipconfig /renew` command will send a query to the DHCP server asking it to resend and renew all DHCP information. For a more detailed look at the `ipconfig` command, type **`ipconfig /?`** at the command prompt. The `ip` command is used with Linux and macOS clients. There is no `/refresh` or `/start` switch for these commands.

15. B. A loopback plug is used to test the ability of a network adapter to send and receive. The plug gets plugged into the NIC, and then a loopback test is performed using troubleshooting software. A toner probe is used to trace a cable. Multimeters test for electrical issues. Cable testers are for verifying that a cable works properly.

16. D. A toner probe is used to trace a cable from one location to another. A loopback plug is for testing the ability of a network adapter to send and receive. Punch-down tools are used to connect the wires of a cable to a wiring rack, such as a 110-block. Cable testers are for verifying that a cable works properly.

17. B. If the hard drive is not recognized, the BIOS/UEFI is the place to check. Most systems do not come with built-in S.M.A.R.T. diagnostics software, even if the hard drive is capable of gathering S.M.A.R.T. data. And if there were one, it would give you potential errors on a working drive. Windows Device Manager is inaccessible without the drive being recognized. There is no boot sector manager, although this issue could be related to the boot sector.

18. A. A crimper is used to make network cables. Punch-down tools are used to connect the wires of a cable to a wiring rack, such as a 110-block. Cable testers are for verifying that a cable works properly (and may be handy after you complete the cable). A loopback plug is for testing the ability of a network adapter to send and receive.

19. C. A RAID 0 array is also known as disk striping. RAID 0 actually decreases your fault tolerance versus one hard drive because there are more points of failure. You need to replace the drive, and hopefully you had it backed up so that you can restore the data.

20. D. The port quickly opening and closing is an example of port flapping. High latency is having a lot of delay on a network, and jitter is variable latency. The port issues could possibly be caused by external interference, but that may not be true as well.

Chapter 13: Mobile Device, Display, and Printer Troubleshooting

1. A, B. Two helpful things to try are toggling the video output function key (usually something like Fn+F8) and plugging an external monitor into the laptop. Removing the display is

possible but not necessary yet. Powering the system off and back on isn't likely to correct the problem.

2. C. Printing preferences sets configuration options such as color vs. black and white, paper size, and page orientation. The print spooler is where print jobs are held for printing. The print queue is the list of print jobs waiting to be printed. There is no paper selection switch on printers (at least not that we've seen).

3. A. With a swollen battery, the best solution is to replace the device. If you are experienced with mobile device batteries, you may be able to replace the battery, but be careful! You should not charge a swollen battery or freeze the system. A battery calibration utility will do no good.

4. B. If print jobs are seemingly getting "stuck" in the printer queue, you should stop and restart the print spooler service. There is no Printer Troubleshooting utility on the physical printer. (There is a Printer Troubleshooter in Windows 11 that might help resolve a print queue issue.) Deleting and reinstalling the printer and/or Windows is not necessary.

5. C, D. Perpetually dim or flickering screens are most likely caused by a failing backlight or inverter, which supplies power to the backlight. LCD devices are not subject to interference from fans. An incorrect video driver will not cause a flickering image.

6. B. The manufacturer's website is the first place you should go for information on your products, including troubleshooting information. Many years ago, manufacturers would provide paper manuals with their products, but that's almost unheard of today, unless you download a PDF version from the website. Server logs can show error codes, but they won't tell you how to fix anything.

7. A. The most likely cause is a groove or a scratch in the EP drum. Toner is collecting in that groove or scratch, and then it is being deposited onto the page. A broken drum-cleaning blade would result in unreadable images (too much toner). Low toner would cause faded images, and a bad transfer corona wire would cause blank images.

8. B. It has to be a problem with the LCD display. If it were the video card, the display would appear warped and fuzzy on the external monitor as well. While many motherboards contain video circuitry, this answer is not specific enough. If the video driver were corrupted, you would have the same problem on all displays.

9. A. If an ink cartridge is faulty or develops a hole, it can release excessive amounts of ink, which will lead to smearing. A corrupted print driver would result in printing garbage. Inkjet printers do not have a fuser. Excessive humidity may cause smearing, but it wouldn't cause the disbursement of too much ink.

10. C, D. Seeing images from previous print jobs is a phenomenon called ghosting (or ghost images). It's most likely due to a bad erasure lamp or a broken cleaning blade. If the transfer corona wire were faulty, there would likely be blank pages. An overheating printer will not necessarily leave residue behind.

11. C. If there is a consistent blank space, it likely means that a pin is not firing properly and the printhead needs to be replaced. If the print ribbon were old, you would have consistently faded printing. If the ribbon were not advancing properly, you would get light and dark printing. If the wrong driver were installed, you would get garbage.

12. D. The first step should be to try to reboot the device. In most cases with a frozen smartphone, holding the power button down for about 10 seconds will force it to power off. If that doesn't work, then you can move to trying to reset the phone by holding down the power button and the sleep/wake button for 10 seconds. There is no need to restore the device to factory settings (yet) or replace the battery.

13. A, B, D. Usually we say to check the easy stuff, and rebooting is almost always a good option. That might work here, but in this case, though, if the tablet seems to be performing normally but the stylus appears to be the problem, focus there. Many styluses have batteries that can go bad, or it's possible the tip has gone bad as well. There is no stylus sensor in tablets. The screen digitizer picks up the signal.

14. C. If the printer is printing garbled output, the most likely cause of the problem is the print driver. Make sure it's the correct driver for the printer. Then, either update the driver, or remove the printer from the OS and reinstall it to get a new driver. It could be the formatter board too, but checking the print driver is far easier. If the printer cable were not seated properly, there likely would be no output. Laser printers do not have printheads.

15. D. If the charger works, what about the outlet? If the phone refuses to charge with either charger or outlet, try restarting the device. If it still won't charge, it could have a problem with the battery or the charging port. An AC versus DC adapter won't make a difference, and there are no secondary auxiliary ports on Android phones.

16. C. If the video works but the audio doesn't, it's probably a settings issue. Either the audio output is set incorrectly, or the volume is turned down or muted. Check those things first. HDMI cables carry audio but DVI cables do not. Rebooting or restarting either device probably isn't necessary if it's a setting that's not configured properly.

17. A. Squealing or grinding noises are made by components that have moving parts. In this case, it's possibly a motor such as the stepper motor. Ink cartridges shouldn't be refilled but they won't make a squealing sound. Paper thickness won't cause squealing, nor will a failing formatter board.

18. D. The symptom here is a display burn-in, also known as a ghost image. The only solution is to replace the monitor.

19. A, B. Paper that is too damp, possibly due to high humidity, can become sticky, making the printer grab more than one sheet at a time. It could also be a tension issue in the paper pickup mechanism. Tray settings affect things like the size of paper or different types of paper such as envelopes. Orientation affects the width and height of output, and finishing issues are more often related to stapling or hole punching.

20. B. The projector is probably shutting off because it's overheating, and that can be caused by dust buildup in the air filter. Check that first, but preferably not when the projector is hot. The mounting bracket likely has nothing to do with it, unless it somehow covers the air intake. Input source will have nothing to do with shutting down and refusing to power back on. If the projector does have auto power-off, it will typically only do that after a certain period of inactivity.

Appendix

B

Answers to Performance-Based Questions

Chapter 1: Motherboards, Processors, and Memory

Answer to Performance-Based Question

Here is how to remove a DIMM and replace it with another one:

1. Pull the tabs on either end of the DIMM away from the DIMM.
2. Pull the loose DIMM straight out of the slot and away from the motherboard.
3. Ensure that the locking tabs are completely opened and out of the way of the slot.
4. Align the module's notch with the tab or tabs in the slot.
5. Insert the new DIMM straight down into the slot.
6. Apply firm and even pressure downward until the locking tabs automatically snap into place.
7. Nudge the tabs inward toward the module to make sure that they are tight.

Answer to Performance-Based Question

The components are labeled in the following illustration.

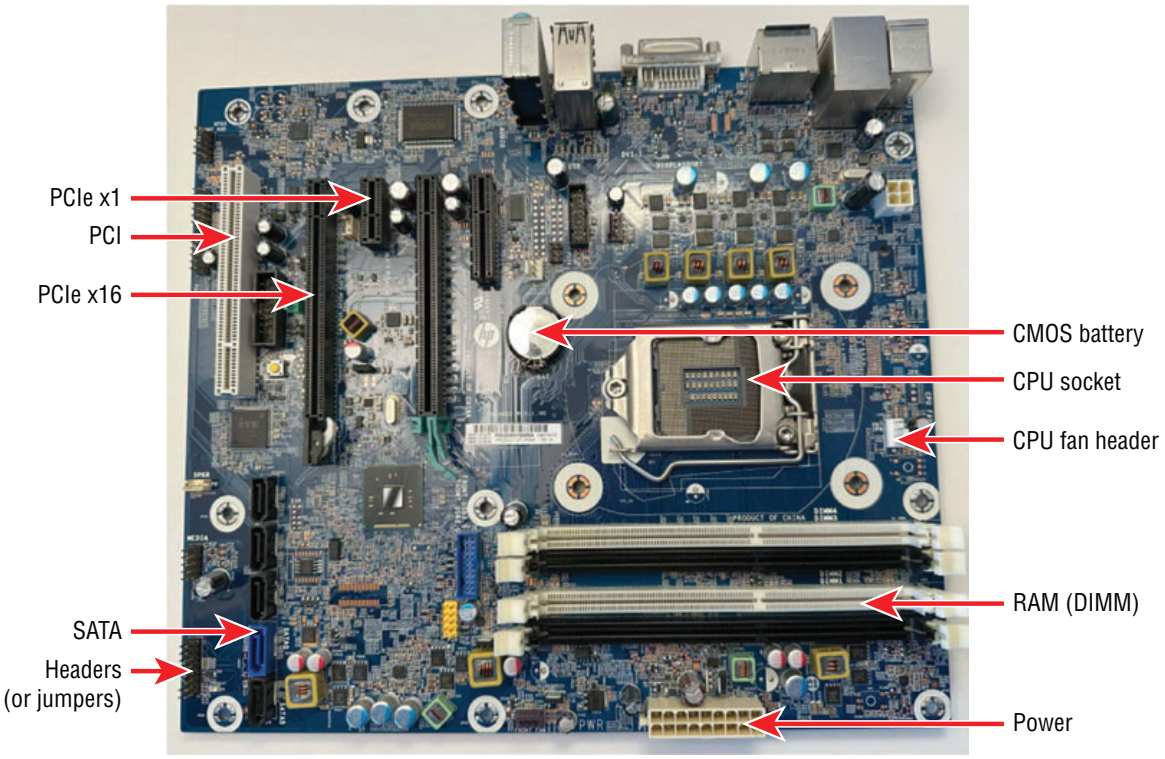

Chapter 2: Expansion Cards, Storage Devices, and Power Supplies

Answer to Performance-Based Question

Here are the steps to remove a power supply from a computer chassis:

1. Remove the power source from the system.
2. Ground yourself and the computer to the same source of ground.
3. Remove the cover from the system.
4. Locate the power supply.
5. Follow all wiring harnesses from the power supply to their termini, disconnecting each one.
6. Remove any obstructions that appear as if they might hinder removal of the power supply.
7. Locate and remove the machine screws on the outside of the case that are used to secure the power supply.
8. Pull the power supply out of the case.

Chapter 3: Peripherals, Cables, and Connectors

Answer to Performance-Based Question

The answers to Chapter 3's, performance-based question are as follows:

- A: VGA
- B: USB
- C: Audio (or 1/8″)
- D: HDMI
- E: eSATA
- F: DisplayPort
- G: USB-C
- H: Mini-DisplayPort

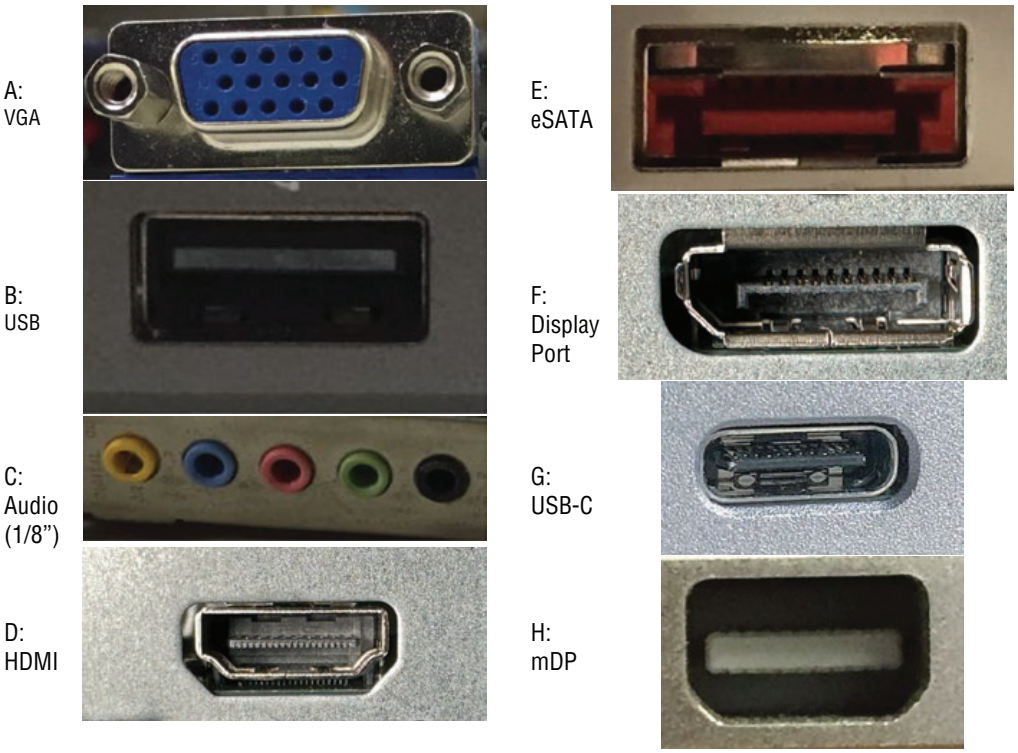

Chapter 4: Laptop and Mobile Device Hardware

Answer to Performance-Based Question

Here is how to replace the hard drive in the example laptop computer.

1. Turn off the computer.
2. Remove the bottom of the case.
3. Remove the screw holding the M.2 drive in place.
4. Slide the hard drive straight out of the M.2 connector.
5. Insert the new drive into the connector at the same angle as the old drive was at when it was unscrewed.
6. Press the drive down and secure it with a screw.
7. Put the bottom of the case back on.

Chapter 5: Printers and Multifunction Devices

Answer to Performance-Based Question

Here are some example steps to take to clean an inkjet printer. The process for starting the cleaning cycle on inkjet printers can vary, and some printers have both quick and deep-clean cycles. Always check your documentation for steps specific to your printer.

1. Power on the printer, and open the top cover to expose the area containing the print cartridges.
2. Initiate a self-cleaning cycle.
3. When the printhead moves from its resting place, pull the AC power plug.
4. Locate the sponge pads on which to apply the cleaning solution.
5. Using the supplied syringe, apply the cleaning solution to the sponge pads until they are saturated.
6. Plug the printer back into the wall outlet and turn it on. The printheads will park themselves.
7. Turn off the printer. Let the solution sit for at least three hours.
8. Power the printer back on and run three printer cleaning cycles. Print a nozzle check pattern (or a test page) after each cleaning cycle to monitor the cleaning progress.

Chapter 6: Networking Fundamentals

Answer to Performance-Based Question

Possible answers for examples of physical network topologies could include bus, ring, star, mesh, and hybrid. The simplest topology, and the one that uses the least amount of cable, is a bus. It consists of a single cable that runs to every workstation, as shown in the following illustration.

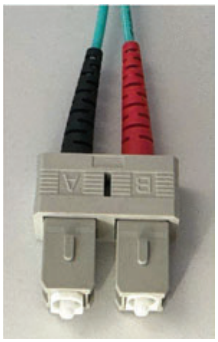
SC

F-type

RJ45 (shielded)

RJ45

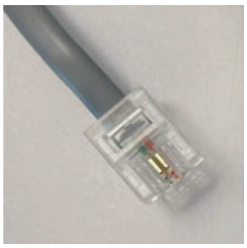
RJ11

ST

LC

BNC

Chapter 7: Introduction to TCP/IP

Answer to Performance-Based Question

Here is the correct matching of protocols and services to their ports:

Protocol (service)	Port(s)
FTP	20–21
SSH	22
Telnet	23
SMTP	25
DNS	53
DHCP	67, 68
HTTP	80
POP3	110
NetBIOS/NetBT	137-139

Protocol (service)	Port(s)
IMAP	143
LDAP	389
HTTPS	443
SMB/CIFS	445
RDP	3389

Chapter 8: Network Services

Answer to Performance-Based Question

The network should look similar to the following illustration.

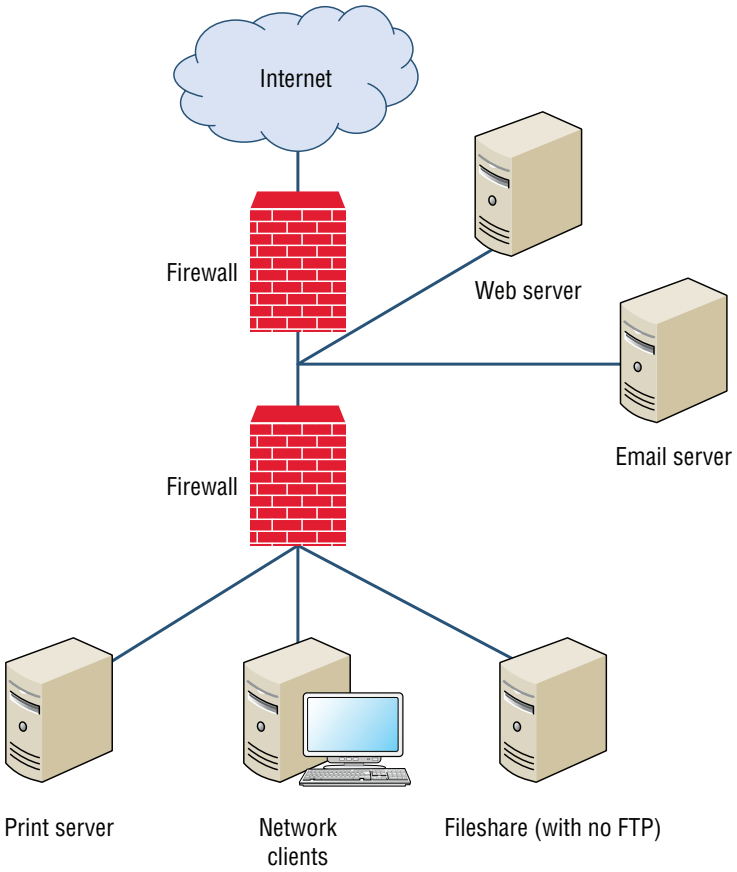

Chapter 9: Wireless and SOHO Networking

Answer to Performance-Based Question

Refer to the figure to see the correct answer.

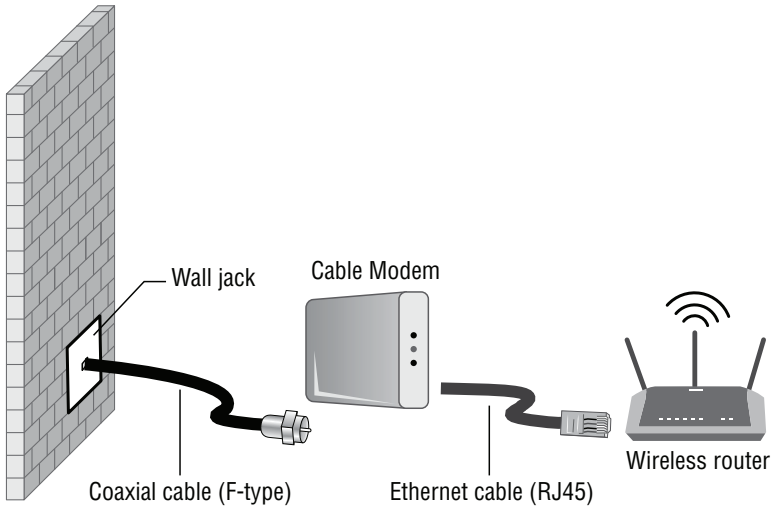

Chapter 10: Mobile Connectivity and Application Support

Answer to Performance-Based Question

Here are the steps to connect an iPhone to a Wi-Fi network:

1. Tap the Settings app on the home screen.
2. Select Wi-Fi from the Settings menu.
3. Swipe the Wi-Fi switch to the right to turn it on if it is off. You can also tap switches to toggle them to the opposite state.

4. In the Choose A Network list, tap the name of the wireless network that you want to join.
5. Enter the password or key for the wireless network, if you are asked for one, and then tap the Join button.

Chapter 11: Virtualization and Cloud Computing

Answer to Performance-Based Question

To enable Microsoft Hyper-V, perform the following steps:

1. Check for minimum system requirements:
   - Windows 11 Enterprise or Pro
   - 64-bit processor with Second Level Address Translation (SLAT)
   - CPU support for VM Monitor Mode Extension (VT-c on Intel CPUs)
   - 4 GB RAM or more
   - Hardware virtualization support enabled in the BIOS/UEFI
2. If requirements are met, then open the Turn Windows Features On Or Off app.
3. Select the Hyper-V box, and click OK.
4. After installation, restart your computer.

Chapter 12: Core Hardware and Network Troubleshooting

Answer to Performance-Based Question

8. Determine the strength of Wi-Fi signals
(Wi-Fi Analyzer)

1. Put connectors on UTP
(Crimper)

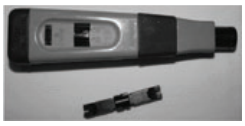

2. Connect wires to a 110-block
(Punchdown tool)

5. Make sure a new cable works
(Cable tester)

7. Trace a cable from end to end
(Toner probe)

6. Test if a NIC is working
(Loopback plug)

3. Cut cables
(Cable stripper)

4. Monitor network traffic
(Network tap)

Chapter 13: Mobile Device, Display, and Printer Troubleshooting

Answer to Performance-Based Question

1. Open the Services app.
2. Find the Print Spooler service.
3. Stop the spooler service. There are several ways that you can do this:
   Right-click the service and choose Stop.
   Click the Stop square above the list of services.
   Click the Stop link to the left of the services.
4. Restart the spooler by right-clicking the service and choosing Start, or by clicking the Start arrow above the list of services.
5. Close the Services app.

Index

A

AAA (authentication, authorization, and accounting), 457–461
AC adapters, as power supplies, 131–132
accelerated graphics port (AGP), 12
access control list (ACL), 379, 550
access points
  about, 372–373
  configuring, 539–553
accessories, setting up, 244–249
accounting, 460–461
ACL (access control list), 379, 550
acronyms, for cellular data connections, 576–577
active hubs, 374
adapter cards
  about, 81, 132
  configuring, 81–94
  configuring devices, 92–94
  exam essentials, 132–133
  exercises, 91–92
  input/output, 88–89
  installing, 81–94
  installing device drivers, 89–92
  multimedia, 84–86
  network interface card (NIC), 86–87
  performance-based question answers, 861
  performance-based questions, 139
  real world scenarios, 87–88
  review question answers, 836–837
  review questions, 134–138
  video, 82–83
address autoconfiguration, 419
Address Resolution Protocol (ARP), 393, 394
ADF (automatic document feeder), 306–307
adjusting display settings, 147–153
ADRAM (asynchronous DRAM), 53–54
ADSL (asymmetric DSL), 514–515
Advanced Micro Devices, Inc. (AMD), 10–11
Advanced RISC Machine (ARM), 45
Advanced tab (Printer Properties), 293
Advanced Technology Extended (ATX)
  about, 7–8
  connectors for motherboards, 123–124
Advanced Television Systems Committee (ATSC), 151
agents, 400–401
AGP (accelerated graphics port), 12
AIO (all-in-one) coolers, 69
air cooling, 66–68
Airplane mode, 571, 573
AirPrint, 302–303
all-black pages, troubleshooting, 814
all-in-one (AIO) coolers, 69
AMD (Advanced Micro Devices, Inc.), 10–11
ammonia, for printers and multifunction devices, 315
Android devices
  configuring location services in, 611–612
  connecting to Wi-Fi networks, 580–582
  disabling cellular use for data networking on, 585–588
  Email configuration on, 619–623
  pairing with Windows laptops, 594–603
  setting up VPNs in, 590–592
  syncing, 631–633
Andy, 659
anycast address, 420–421
Anything/everything as a service (XaaS), 670
APIPA (Automatic Private IP Addressing), 418–419, 759
Application layer, in Open Systems Interconnection (OSI) model, 345
application virtualization, 652
applications, verifying compatibility with, 295–296
ARC (Audio Return Channel), 178

architecture, of central processing unit (CPU), 43–46
ARM (Advanced RISC Machine), 45
ARP (Address Resolution Protocol), 393, 394
aspect ratio, 150–151
assembly, for laser scanning, 271–272
asymmetric DSL (ADSL), 514–515
asynchronous DRAM (ADRAM), 53–54
ATSC (Advanced Television Systems Committee), 151
ATX (Advanced Technology Extended)
  about, 7–8
  connectors for motherboards, 123–124
audio, troubleshooting, 781–782
audio devices, 157–158
audio jacks, on motherboard headers, 31
Audio Return Channel (ARC), 178
audit logs, 304–306
authentication, authorization, and accounting (AAA), 457–461
authentication, troubleshooting failures, 763
authorization, 460
automatic document feeder (ADF), 306–307
Automatic Private IP Addressing (APIPA), 418–419, 759
auxiliary devices, 378–381
availability, cloud computing and, 673–674

B

backlight, 144, 208
back-off time, 490–491
badging, 303
bare-metal hypervisor, 652
baseband OS, 574
Basic Input/Output System (BIOS)
  about, 31–33
  configuring, 33–34
  exercises, 34–36
  flashing the system, 39–40, 711
  monitoring, 33–34
  password, 36
  power-on self-test (POST) routine, 39
  security and encryption, 36–39
  troubleshooting, 710–712
Basic Rate (BR), 506
Basic Service Set (BSS), 501–502
basis weight, of printer paper, 310
batteries
  battery backup systems, 128–130
  CMOS, 40–41
  for laptops, 233–235
  troubleshooting, 784–786
BD (Blu-ray Disc) drive, 115
BEDO RAM (burst EDO DRAM), 53–54
beep code, 713
bias voltage, 278
bidirectional, 14
Binary Phase Shift Keying (BPSK), 494
binary values, 410
biometrics, for laptops and mobile devices, 241–242
BIOS (Basic Input/Output System)
  about, 31–33
  configuring, 33–34
  exercises, 34–36
  flashing the system, 39–40, 711
  monitoring, 33–34
  password, 36
  power-on self-test (POST) routine, 39
  security and encryption, 36–39
  troubleshooting, 710–712
black lines, troubleshooting in printing, 815
blank pages, troubleshooting, 812–813
blank screen, troubleshooting, 715
blanks, 707
Blue Screen of Death (BSOD), 716
Bluetooth
  about, 329–330, 505
  devices, 507–508
  establishing connectivity, 593–606
  for laptops and mobile devices, 245
  for printers and multifunction devices, 284–285
  versions, 505–507
Blu-ray Disc (BD) drive, 115
BNC connector, 356
Bonjour, 301–302
boot options, 33
boot password, 36
BPaaS (business processes as a service), 670
BPSK (Binary Phase Shift Keying), 494
BR (Basic Rate), 506

bring your own device (BYOD), 614
broadcasting, 373
BSOD (Blue Screen of Death), 716
BSS (Basic Service Set), 501–502
Buhagiar, Jon (author)
CompTIA Network+ Study Guide, 325
burn-in, 146
burning smell, troubleshooting, 708–709
burnt-out bulb, 776
burst EDO DRAM (BEDO DRAM), 53–54
bus architecture, 9–10
bus topology, 341
business applications, syncing, 633–637
business processes as a service (BPaaS), 670
BYOD (bring your own device), 614

C

CA (certificate authority), 402
CaaS (communications as a service), 670
cable Internet, 518–519
cable locks, for laptops and mobile devices, 242–243
cable modems
about, 372
installing and configuring, 537
cable strippers, 737
cable tester, 739, 740
cables and connectors
about, 143, 162–163, 184, 353–371
exam essentials, 185
hard drive, 181–184
for motherboards, 11–15
for networking, 353–371
performance-based question, 190
performance-based question answers, 861–862
peripheral, 163–172
for power supplies, 123–126
review question answers, 838–839
review questions, 186–189
serial ports, 172–173
troubleshooting, 714–715, 736–741
video, 173–181
cache memory, 16–18
calendars, syncing, 633–637
calibrating
about, 267
printers, 288
caliper, of printer paper, 311
Canonical Name (CNAME), 444
capacitative touchscreens, 146
capacitor swelling, troubleshooting, 715–716
carriage belt, 263
carriage motor, 263
carriage stepper motor, 263
carrier, 270
Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA), 348, 490–491
Carrier Sense Multiple Access with Collision Detection (CSMA/CD), 347–348
case, laptop, 202–206
case cooling, 63–64
categories, for twisted pair cables, 360–361
cathode ray tube (CRT), 146
CD (compact disc), 115
CDMA (code-division multiple access), 565
CEC (Consumer Electronics Council), 177
Celeron, 21–22
cellular data connections
about, 567
acronyms, 576–577
cellular updates, 574–576
enabling and disabling wireless/cellular connections, 570–573
mobile hotspots, 567–570
tethering, 567–570
cellular Internet, 523–525
cellular location services, 608–609
cellular networking standards
about, 565
fifth generation (5G), 524–525, 566–567
fourth generation (4G), 523–524, 566
third generation (3G), 565
cellular updates, 574–576
central processing unit (CPU)
about, 18–21
architecture of, 43–46
characteristics of, 46–49
cooling, 66–69
centralized processing, 327

certificate authority (CA), 402
CF (CompactFlash), 113
channels
  changing, 495–496
  configuring, 52–53
  Wi-Fi and, 491–492
character set, 345
characteristics, CPU, 46–49
characters per second (cps), 259
charging corona, 276
charging step, in electrophotographic (EP) imaging process, 276, 808
chip creep, 707–708
chipsets, 10–11
Chromebooks, 194
CIDR (classless inter-domain routing), 414–416
CIFS (Common Internet File System), 402
CISC (Complex Instruction Set Computing), 43–44
clamshell design, 203
Class A networks, 412
Class B networks, 413
Class C networks, 413
classless inter-domain routing (CIDR), 414–416
cleaning solutions, for printers, 313–315
cleaning step, in electrophotographic (EP) imaging process, 279–280, 808
client computers, 334
client software, 334–335
client-server resource model, 338–340
client-side virtualization
  about, 655
  requirements
    about, 656
    emulator requirements, 659
    resource requirements, 656–659
    security requirements, 659–667
cloud computing
  about, 649, 668, 688–689
  characteristics of clouds, 671–674
  cloud services, 669–671
  cloud-based services
    cloud-based applications, 675–678
    cloud-based storage, 674–675
    setting up cloud servers, 678–688
  exam essentials, 689–690
  exercises, 676–677, 678–688
  performance-based question, 695
  performance-based question answers, 866–867
  review question answers, 853–854
  review questions, 691–694
  types of clouds, 668–669
cloud servers, setting up, 678–688
cloud storage, 637–639
cloud-based applications, 675–678
cloud-based services
  about, 669–671
  cloud-based applications, 675–678
  cloud-based storage, 674–675
  setting up cloud servers, 678–688
cloud-based storage, 674–675
CMOS (complementary metal oxide semiconductor), 40–41
CMOS battery, 711
CMR (conventional magnetic recording), 98–99
CNAME (Canonical Name), 444
coaxial cable, for networking, 354–358
code-division multiple access (CDMA), 565
coding scheme, Wi-Fi and, 492
color gamut, 155
color output problems, troubleshooting, 805
Common Internet File System (CIFS), 402
communication
  interfaces for SSDs, 100–105
  for networking, 344–348
communications as a service (CaaS), 670
community cloud, 669
compact disc (CD), 115
CompactFlash (CF), 113
compatibility
  network interface card (NIC) and, 349
  verifying, 295–296
complementary metal oxide semiconductor (CMOS), 40–41
Complex Instruction Set Computing (CISC), 43–44
composition, of printer paper, 310
CompTIA Network+ Study Guide (Lammle and Buhagiar), 325
configuring

access points, 539–553
devices, 614
Email accounts, 615–626
expansion cards, 81–94
laptop hardware, 202–244
location services, 610–612
mobile device synchronization, 626–639
network interface card (NIC), 351–352
printer options/settings, 289–296
VPN clients, 427–429
wireless firewalls, 551
wireless routers, 539–553
connectivity
devices
about, 371–378
installing and configuring, 536–539
for networking, 371–378
troubleshooting, 756–763
methods for laptops and mobile devices, 244–245
mobile
about, 563–565, 639
cellular data connections, 567–577
cellular networking standards, 565–567
establishing Bluetooth connectivity, 593–606
establishing Wi-Fi connectivity, 577–593
exam essentials, 639–640
exercises, 578–588, 590–592, 594–603
performance-based question, 645
performance-based question answers, 866–867
real world scenarios, 567–568, 574
review question answers, 851–852
review questions, 641–644
troubleshooting, 788–792
connectors and cables
about, 143, 162–163, 184, 353–371
exam essentials, 185
hard drive, 181–184
for motherboards, 11–15
for networking, 353–371
performance-based question, 190
performance-based question answers, 861–862
peripheral, 163–172
for power supplies, 123–126
review question answers, 838–839
review questions, 186–189
serial ports, 172–173
troubleshooting, 714–715, 736–741
video, 173–181
consumables, replacing, 310–313
Consumer Electronics Council (CEC), 177
contacts, syncing, 633–637
container engine, 655
container image, 655
containers, 655–656
contention-based, 348
controlling data, network interface card (NIC) and, 349–350
conventional magnetic recording (CMR), 98–99
cooling systems
about, 61–63, 71
case cooling and fans, 63–64
CPU cooling, 66–68
exam essentials, 71–72
hard drive cooling, 70
liquid cooling, 69
memory cooling, 70
motherboard fan power connectors, 64–65
performance-based question answers, 860
performance-based questions, 77
review question answers, 834–835
review questions, 73–76
video card cooling, 70
cores, CPU, 46
corona wire, 274
corporate applications, 613
corporate devices, 614
cost
of inkjet printers, 261
laptops compared with desktops, 195
cps (characters per second), 259
CPU (central processing unit)
about, 18–21
architecture of, 43–46
characteristics of, 46–49
cooling, 66–69
CPU-Z, 18, 48–49

crossover cable, 363
cross-platform virtualization, 652
CRT (cathode ray tube), 146
CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance), 348, 490–491
CSMA/CD (Carrier Sense Multiple Access with Collision Detection), 347–348
cursor drift, 788

D

daisy-wheel printers, 258–259
damage, troubleshooting, 708–709
data as a service (DaaS), 670
data availability, of database servers, 464
data caps, 627–628
data in transit, 402
Data Link layer, in Open Systems Interconnection (OSI) model, 346
data migration, in laptops, 224
data persistence, of database servers, 464
data rate, Wi-Fi and, 492–493
data variety, of database servers, 464
database management system (DBMS), 461–462
database servers, 461–465
databases, types of, 464–465
daughterboard, 212–213
DB-9 serial connector, 172–173
DBMS (database management system), 461–462
DC power supply (DCPS), for laser printers, 272
dead pixels, 779–780
dedicated servers, 333, 440
default allow, 550
default deny, 379, 550
default gateway, 412
defects, troubleshooting in printing, 814–815
degraded performance, 797–798
demilitarized zone (DMZ), 379, 441
desktop as a service (DaaS), 670
desktops, laptops compared with, 194–196
developer, 270
developing roller, 278
developing step, in electrophotographic (EP) imaging process, 277–278, 808
device drivers, installing and updating, 89–92
Device Manager, 89–92
devices
  Android
    configuring location services in, 611–612
    connecting to Wi-Fi networks, 580–582
    disabling cellular use for data networking on, 585–588
    Email configuration on, 619–623
    pairing with Windows laptops, 594–603
    setting up VPNs in, 590–592
    syncing, 631–633
  audio, 157–158
  auxiliary, 378–381
  Bluetooth, 507–508
  configuring, 614
  802.11, 504–505
  hot-swappable, 113–114
  input/output
    about, 159
    keyboard, 159–160
    KVM switches, 161–162
    mouse, 160–161
  intermittent failure of, 710
  Internet of Things (IoS)
    about, 475
    home security and automation, 477–479
    thermostats, 476
    voice-enabled smart speakers/digital assistants, 479
  iPhones
    connecting to Wi-Fi networks, 578–580
    disabling cellular use for data networking on, 582–585
    Email configuration on, 616–619
    pairing, 603–605
  storage
    about, 81, 94, 132
    exam essentials, 132–133
    exercises, 117–119
    hard disk drive (HDD) systems, 94–99

installing, 116, 118–119
performance-based question answers, 861
performance-based questions, 139
as peripherals, 162
Redundant Array of Independent Disks (RAID), 109–111
removable, 111–116
removing, 116–118
review question answers, 836–837
review questions, 134–138
solid-state drives (SSDs), 99–108
video
about, 144
monitors, 144–155
projection systems, 155–157
"devices not found," 724–725
DFS (dynamic frequency selection), 497, 500
DHCP (Dynamic Host Configuration Protocol)
about, 399, 448
how it works, 450
scopes and reservations, 448–449
DHCP relay agent, 450
DHCP server, 448
digital subscriber line (DSL)
about, 514–517
modem, 372, 537
digital versatile disc (DVD), 115
Digital Visual Interface (DVI), 175–177
digitizer, 146, 208
dim image, 781
DIMM (dual in-line memory module), 16, 59–61
DIP (dual in-line package), 42
direct burial cable, 364–365
direct-sequence spread spectrum (DSSS), 493
disassembling laptops, 199–202
disk mirroring (RAID 1), 109, 111
disk striping (RAID 0), 109
display, for laptops, 209–212
display burn-in, 779
DisplayPort, 180–181
displays
laptop, 206–212
troubleshooting, 781–783
distended capacitors, 715–716
distorted image, 778–779
distributed processing, 327
DKIM (Domain Keys Identified Mail), 445
DMARC (Domain-based Message Authentication, Reporting, and Conformance), 445
DMZ (demilitarized zone), 379, 441
DNAT (Dynamic Network Address Translation), 552
DNS (Domain Name System)
about, 398–399, 442–443
DNS servers, 443–444
on the Internet, 445–447
spam management, 444–445
docking stations, 247–249
domain controller, 339, 458
Domain Keys Identified Mail (DKIM), 445
Domain Name System (DNS)
about, 398–399, 442–443
DNS servers, 443–444
on the Internet, 445–447
spam management, 444–445
Domain-based Message Authentication, Reporting, and Conformance (DMARC), 445
domains, 339
door locks, 477–478
dot-matrix printers, 259–260
double data rate SDRAM, 54, 55–58
double-parity RAID (RAID 6), 110, 111
draft quality, 260
DRAM (dynamic random access memory), 53–58
drawing pad, 247
drive interfaces, 182
driver software, updating, 286–288
drivers, network interface card (NIC), 352–353
drives, troubleshooting, 722–735
DSL (digital subscriber line)
about, 514–517
modem, 372, 537
DSSS (direct-sequence spread spectrum), 493
D-subminiature connector, 174–175

dual in-line memory module (DIMM), 16, 59–61
dual in-line package (DIP), 42
dual-channel memory, 52
dual-rail power supplies, 121–122
dual-voltage options, 121
duplex printing, 294
duplexing, 109
duplexing assembly
  about, 266
  for laser printers, 275
DVD (digital versatile disc), 115
DVI (Digital Visual Interface), 175–177
DVI-to-HDMI adapter, 178–179
DVI-to-VGA adapter, 176
dynamic frequency selection (DFS), 497, 500
Dynamic Host Configuration Protocol (DHCP)
  about, 399, 448
  how it works, 450
  scopes and reservations, 448–449
dynamic IP address, 403, 449
Dynamic Network Address Translation (DNAT), 552
dynamic random access memory (DRAM), 53–58

E

ECC (error connection code), 51
EDO DRAM (extended data out DRAM), 53–54
EDR (Enhanced Data Rate), 506
EEPROM (electronically erasable PROM), 59
802.11 networking
  about, 490–491
  devices, 504–505
  gigabit standards, 499–504
  legacy standards, 492–499
  terminology for, 491–492
802.11a standard, 493
802.11ac standard, 499–500
802.11ax standard, 500–502
802.11b standard, 495
802.11be standard, 502–503
802.11g standard, 495
802.11n standard, 497–498
80 PLUS certification, 123
elasticity, 671
electromagnetic interference (EMI), 528
electronic stepper motor, 273
electronically erasable PROM (EEPROM), 59
electrophotographic (EP) imaging process, 275–281, 807–808
Email accounts
  configuring, 615–626
  syncing, 634–637
eMBB (Enhanced Mobile Broadband), 566
embedded systems, 473–475
EMI (electromagnetic interference), 528
emulator requirements, for virtualization, 659
enabling and disabling wireless/cellular connections, 570–573
encryption, BIOS/UEFI and, 36–39
energy efficiency, 122–123
Enhanced Data Rate (EDR), 506
Enhanced Mobile Broadband (eMBB), 566
environment, for printers and multifunction devices, 315
EP (electrophotographic) imaging process, 275–281, 807–808
erasable PROM (EPROM), 59
ergonomic keyboard, 160
error checking, 51
error connection code (ECC), 51
eSATA (external SATA)
  about, 27
  connectors, 182–183
eSATA expansion cards, 88–89
establishing
  Bluetooth connectivity, 593–606
  Wi-Fi connectivity, 577–593
Ethernet
  about, 10
  naming standards for, 355
  for printers and multifunction devices, 284
exam essentials
  cables and connectors, 185
  cloud computing, 689–690
  cooling systems, 71–72
  expansion cards, 132–133

laptop and mobile device hardware, 249
memory, 71–72
mobile application support, 639–640
mobile connectivity, 639–640
motherboards, 71–72
network services, 480–481
networking, 382–383
peripherals, 185
power supplies, 132–133
printers and multifunction devices, 316–317
processors, 71–72
small office/home office (SOHO) networks, 554–555
storage devices, 132–133
Transmission Control Protocol/Internet Protocol (TCP/IP), 430–431
troubleshooting, 763–765, 824–826
virtualization, 689–690
wireless networks, 554–555
exclusion, 448
exercises
Basic Input/Output System (BIOS), 34–36
cloud computing, 676–677, 678–688
expansion cards, 91–92
laptop and mobile device hardware, 204–206, 210–212, 214–216, 218–219, 222, 225–226, 227–228, 229–230, 234, 236–241
mobile application support, 610–612, 616–623
mobile connectivity, 578–588, 590–592, 594–603
network services, 470–472
networking, 371
peripherals, 147–149, 151–153
power supplies, 131
printers and multifunction devices, 267, 297, 299–300, 301–302, 308–309, 314
small office/home office (SOHO) networks, 529–531
storage devices, 117–119
Transmission Control Protocol/Internet Protocol (TCP/IP), 403–409, 427–429
troubleshooting, 721–722, 726–729, 743–745, 751–753, 820–821
Unified Extensible Firmware Interface (UEFI), 34–36
virtualization, 653–654, 661–667
expandability, laptops compared with desktops, 195
expansion cards
about, 81, 132
configuring, 81–94
configuring devices, 92–94
exam essentials, 132–133
exercise, 91–92
input/output, 88–89
installing, 81–94
installing device drivers, 89–92
multimedia, 84–86
network interface card (NIC), 86–87
performance-based question answers, 861
performance-based questions, 139
real world scenario, 87–88
review question answers, 836–837
review questions, 134–138
video, 82–83
expansion slots, 11–15
exposing step, in electrophotographic (EP) imaging process, 276–277, 808
extended data out DRAM (EDO DRAM), 53–54
extenders, 373
external SATA (eSATA)
about, 27
connectors, 182–183
eyestrain, 149

F

facility code, 455
fan considerations, 34
fans, 63–64
fast page mode DRAM (FPM DRAM), 53–54
feed roller, 273
feeder, 263–266
FHSS (frequency-hopping spread spectrum), 493

fiber-optic cables, for networking, 365–371
fiber-optic Internet, 519–520
fiber-to-the-curb (FTTC), 519
fiber-to-the-home (FTTH), 519
fiber-to-the-node (FTTN), 519
fifth generation (5G), 524–525, 566–567
file locking, 328
file server, 333, 451
file synchronization, cloud computing and, 674
File Transfer Protocol (FTP), 393, 396–397
fileshare, 451
firewalls, 378–379, 549–551
firmware
  about, 31
  updating, 288
fixed cable PSU, 126
flash memory, 111–113
flashing screen, 780
flashing the system BIOS/UEFI, 39–40, 711
flat file, 463
flatbed scanner, 306
Fn key, 787
form factors
  for HDD systems, 98
  motherboard, 6–8
  SSDs, 105–108
formatter board, 817
fourth generation (4G), 523–524, 566
FPM DRAM (fast page mode DRAM), 53–54
FQDN (fully qualified domain name), 446
frames, 346
frequency, Wi-Fi and, 491
frequency-hopping spread spectrum (FHSS), 493
front intake fan, 63
frontend, 327
FSB speed, 46–47
FTP (File Transfer Protocol), 393, 396–397
FTTC (fiber-to-the-curb), 519
FTTH (fiber-to-the-home), 519
FTTN (fiber-to-the-node), 519
F-type connector, 357, 358
full-duplex communication, 350
fully qualified domain name (FQDN), 446
fusing assembly, for laser printers, 274
fusing step, in electrophotographic (EP) imaging process, 279, 808
fuzzy image, 776–777

G

General tab (Printer Properties), 290, 291
ghost cursor, 787
ghosting, troubleshooting in printing, 816
gigabit 802.11 standards, 499–504
Global Navigation Satellite Systems (GNSSs), 607–608
Global Positioning System (GPS), 606–608
Global System for Mobile Communications (GSM), 565
Google Drive, scanning documents to, 308–309
graphics processing unit (GPU), 82
grinding, troubleshooting, 705–706

H

HaaS (hardware as a service), 670
half-duplex communication, 350
hard disk drive (HDD) systems
  about, 94–95
  anatomy of, 96
  form factors, 98
  spindle speeds, 96–98
hard drives
  cable and connectors, 181–184
  cooling, 70
  for laptops, 220–224
hardware. *See also* troubleshooting
  networking
    about, 348
    auxiliary devices, 378–381
    cables and connectors, 353–371
    connectivity devices, 371–378
    network interface card (NIC), 348–353
  troubleshooting tools for, 736–741
hardware as a service (HaaS), 670
hardware security module (HSM), 37
HDD (hard disk drive) systems

about, 94–95
anatomy of, 96
form factors, 98
spindle speeds, 96–98
HDMI (High-Definition Multimedia Interface), 177–180
HDMI-to-VGA adapter, 179
HDSL (high bit-rate DSL), 514–515
headaches, 149
headers
motherboard, 28–31
in Open Systems Interconnection (OSI) model, 345
headset, 246
heat
for printers and multifunction devices, 315
troubleshooting issues, 784–786
heat sinks, 63
heat spreaders, 70
hierarchical addressing scheme, 409
high bit-rate DSL (HDSL), 514–515
High Speed (HS), 506
High-Definition Multimedia Interface (HDMI), 177–180
high-voltage power supply (HVPS)
about, 813
for laser printers, 272
Hi-Speed, 164–165
home security and automation, 477–479
host, 403
host ID, 411
host-based firewalls, 550
`hostname` command, 746–747
Host-to-Host layer protocols, 394–396
hotspot, for laptops and mobile devices, 245
hot-swappable devices, 113–114
HS (High Speed), 506
HSM (hardware security module), 37
HTT (Hyper-Threading Technology), 46–48
HTTP (Hypertext Transfer Protocol), 393, 399
HTTPS (Hypertext Transfer Protocol Secure), 393, 401–402, 456
hub-and-spoke topology, 341–342
hubs
about, 373–375
installing and configuring, 537
humidity, for printers and multifunction devices, 315
HVPS (high-voltage power supply)
about, 813
for laser printers, 272
hybrid cloud, 669
hybrid topology, 344
Hypertext Transfer Protocol (HTTP), 393, 399
Hypertext Transfer Protocol Secure (HTTPS), 393, 401–402, 456
Hyper-Threading Technology (HTT), 46–48
hypervisor, 652–655

I

IaaS (infrastructure as a service), 669–670
ICCID (integrated circuit card identifier), 576
iCloud, 628, 629, 630
ICMP (Internet Control Message Protocol), 393, 394
ICSs (industrial control systems), 475
IDE (integrated drive electronics) drives, 25–26
IDS (intrusion detection system), 467
IEEE (Institute of Electrical and Electronics Engineers), 329, 490
IEEE 802 standards, 346–348
IEEE 1394, 175
IMAP (Internet Message Access Protocol), 393, 400, 454
IMEI (International Mobile Equipment Identity), 576
impact paper, 259
impact printers
about, 258
daisy-wheel, 258–259
dot-matrix, 259–260
troubleshooting, 800–803
IMSI (international mobile subscriber identity), 576
incorrect color display, 781

industrial control systems (ICSs), 475
Information Technology Extended (ITX), 8
infrastructure, 330
infrastructure as a service (IaaS), 669–670
ink cartridge, 261–262, 311–313
inkjet printers
  about, 260
  components of, 260–267
  cost of, 261
  printing process, 267–269
  troubleshooting, 803–807
innovation, speed of, 14
in-plane switching (IPS), 145
input
  devices for laptops, 225–231
  methods for mobile devices, 197–198
  troubleshooting issues, 786–788
  using displays for, 146–147
input/output card
  about, 88–89
  power supply, 120–121
input/output devices
  about, 159
  keyboard, 159–160
  KVM switches, 161–162
  mouse, 160–161
input/output operations per second (IOPS), 729
installing
  device drivers, 89–92
  expansion cards, 81–94
  laptop hardware, 202–244
  network infrastructure
    about, 530
    connectivity devices, 536–539
    internal NIC, 530–535
    network cables, 535–536
  networked printers, 298–309
  printers and multifunction devices, 282–315
  storage devices, 116, 118–119
Institute of Electrical and Electronics Engineers (IEEE), 329, 490
integrated circuit card identifier (ICCID), 576
integrated drive electronics (IDE) drives, 25–26
integrated print server, 299
interface circuitry, 267
intermittent device failure, 710
internal expansion, for laptops, 231–233
internal network connections, choosing, 526–530
internal NIC, installing, 530–535
International Mobile Equipment Identity (IMEI), 576
international mobile subscriber identity (IMSI), 576
Internet appliances
  about, 465
  load balancing, 468–470
  proxy server, 470–473
  spam gateways, 465–466
  Unified Threat Management (UTM), 466–467
Internet connections, choosing, 514–526
Internet Control Message Protocol (ICMP), 393, 394
Internet Key Exchange v2 (IPEv2), 590
Internet layer protocols, 394
Internet Message Access Protocol (IMAP), 393, 400, 454
Internet of Things (IoS) devices
  about, 475
  home security and automation, 477–479
  thermostats, 476
  voice-enabled smart speakers/digital assistants, 479
Internet Protocol (IP), 394
Internet Protocol Security (IPSec), 426, 590, 592–593
Internet service provider (ISP), 514
Internetwork Operating System (IOS), 112
intrusion detection system (IDS), 467
Intrusion prevention system (IPS), 467
inverter, 144, 208
I/O board, 212–213
I/O ports, troubleshooting, 714–715
Iometer, 729
IOPS (input/output operations per second), 729
iOS

configuring location services in, 610–611
syncing devices, 628–631
IOS (Internetwork Operating System), 112
IoS (Internet of Things) devices
about, 475
home security and automation, 477–479
thermostats, 476
voice-enabled smart speakers/digital assistants, 479
IP (Internet Protocol), 394
IP addressing
about, 403–409
Automatic Private IP Addressing (APIPA), 418–419
binary values, 410
classless inter-domain routing (CIDR), 414–416
components of, 411–412
conflicts with, 759–760
IPv4 address classes, 412–414
IPv6, 420–422
public vs private, 416–418
`ipconfig` command, 741–745
IPEv2 (Internet Key Exchange v2), 590
iPhones
connecting to Wi-Fi networks, 578–580
disabling cellular use for data networking on, 582–585
Email configuration on, 616–619
pairing, 603–605
IPS (in-plane switching), 145
IPS (Intrusion prevention system), 467
IPSec (Internet Protocol Security), 426, 590, 592–593
IPv4 address classes, 412–414
IPv4 host record (A), 444
IPv6, 420–422
ISP (Internet service provider), 514
iTunes, 628–631
ITX (Information Technology Extended), 8

J

jailbreaking, 795
jitter, 760

K

Kerberos, 458
keyboards
about, 159–160
for laptops, 225–226
KVM switches, 161–162

L

L1 cache, 16
L2 cache, 17
L2TP (Layer 2 Tunneling Protocol), 590
L3 cache, 17
Lammle, Todd (author)
CompTIA Network+ Study Guide, 325
land grid array (LGA), 19
lands, on CPUs, 42–43
lanes, for PCIe, 13
LANs (local area networks), 327–328
Laplink PCmover, 224
laptop and mobile device hardware
about, 193–194, 249
characteristics of mobile devices, 196–199
configuring, 202–244
configuring accessories and ports, 244–249
disassembling laptops, 199–202
exam essentials, 249
exercises, 204–206, 210–212, 214–216, 218–219, 222, 225–226, 227–228, 229–230, 234, 236–241
installing, 202–244
laptops compared with desktops, 194–196
malware on laptops, 794–795
performance-based question, 254
performance-based question answers, 862
real world scenarios, 207–208, 235
reassembling laptops, 199–202
review question answers, 839–841
review questions, 250–253
setting up accessories and ports, 244–249
laser printers
about, 269–270
components of, 270–275

electrophotographic (EP) imaging process, 275–281
troubleshooting, 807–817
latency, 521, 760
Layer 2 Tunneling Protocol (L2TP), 590
LCD (liquid crystal display), 144, 146, 194
LCD cutoff switch, 775
LDAP (Lightweight Directory Access Protocol), 401
LE (Low Energy), 506
LEDs (light-emitting diodes), 145–146, 723
legacy 802.11 standards, 493–499
legacy systems, 473–475
letter quality (LQ), 259
LGA (land grid array), 19
light, for printers and multifunction devices, 315
light switches, 478–479
light-emitting diodes (LEDs), 145–146, 723
Lightning connector
about, 169–170
for laptops and mobile devices, 245
lights, troubleshooting, 723–724
Lightweight Directory Access Protocol (LDAP), 401
line-of-sight, 522
link local address, 422
links, for PCIe, 13
liquid cooling, 69
liquid crystal display (LCD), 144, 146, 194
lm (lumen), 156
load balancing, 468–470
local area networks (LANs), 327–328
local connector, 370
local printers, installing and sharing, 282–298
location services
about, 606
cellular, 608–609
configuring, 610–612
Global Positioning System (GPS), 606–608
Long-Term Evolution (LTE), 524
loopback plug, 739, 740
Low Energy (LE), 506
low-resolution mode, 782
LQ (letter quality), 259
LTE (Long-Term Evolution), 524
lumen (lm), 156

M

M.2, 28, 106–108, 232–233
MAC address (Media Access Control address), 346, 351–352
MAC Filtering, 546
Macrium Reflect X, 224
Mail Exchange (MX), 444
mail server, 452–454
mainboard. *See* motherboards
maintaining printers and multifunction devices, 282–315
maintenance kits, 808
maintenance station, 262
malicious apps, 795
malware, troubleshooting, 792, 794–796
MAM (mobile application management), 613–626
managed switch, 376
managing
print jobs, 817–819
print options, 821–823
print spooler, 820–821
MANs (metropolitan area networks), 330
manufacturer documentation, 199–200
marks, troubleshooting in printing, 814–815
Massive Machine Type Communications (mMTC), 566
MCC (mobile country code), 576
MDM (mobile device management), 613–626
mDP (Mini DisplayPort) connector, 180
measured service, 671–672
Media Access Control address (MAC address), 346, 351–352
MEID (Mobile Equipment Identifier), 576
memory
about, 5–6, 50, 71
cooling, 70
exam essentials, 71–72
for laptops, 217–220
packaging, 59–61

performance-based question answers, 860
performance-based questions, 77
review question answers, 834–835
review questions, 73–76
terminology for, 50–53
troubleshooting, 716–719
types of, 53–59
memory slots, 16–18
mesh topology, 343–344
metered utilization, 671–672
metropolitan area networks (MANs), 330
MFA (multifactor authentication), 458
microATX, 8
MicroDIMM, 218
microphone, 209
microSD, 113
Microsoft 365, 636–637
Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2), 592–593
Microsoft's BitLocker, 37
micro-USB connector, 168
MIDI connections, 85
migration software, 224
Mini DisplayPort (mDP) connector, 180
Mini light-emitting diode (Mini-LED), 145
Mini PCIe, 231–232
mini-HDMI connector, 179–180
mini-ITX, 8
miniSD, 113
mini-Serial ATA (mSATA), 105–106
mini-USB connector, 168
MLO (Multi-Link Operation), 502
MMF (multimode fiber), 366
mMTC (Massive Machine Type Communications), 566
mmWave, 525, 567
MNC (mobile network code), 576
mobile application management (MAM), 613–626
mobile application support
about, 563, 606, 639
configuring mobile device synchronization
about, 626–627
cloud storage, 637–639
data caps, 627–628
syncing Android devices, 631–633
syncing Apple iOS devices, 628–631
syncing calendars, contacts, and business applications, 633–637
exam essentials, 639–640
exercises, 610–612, 616–623
location services
about, 606
cellular, 608–609
configuring, 610–612
Global Positioning System (GPS), 606–608
mobile device and app management
about, 613–614
configurations, 614
configuring Email accounts, 615–626
MDM policy enforcement, 614–615
performance-based question, 645
review question answers, 851–852
review questions, 641–644
mobile connectivity
about, 563–565, 639
cellular data connections
about, 567
acronyms, 576–577
cellular updates, 574–576
enabling and disabling wireless/cellular connections, 570–573
mobile hotspots, 567–570
tethering, 567–570
cellular networking standards
about, 565
fifth generation (5G), 566–567
fourth generation (4G), 566
third generation (3G), 565
establishing Bluetooth connectivity, 593–606
establishing Wi-Fi connectivity, 577–593
exam essentials, 639–640
exercises, 578–588, 590–592, 594–603
performance-based question, 645
performance-based question answers, 866–867
real world scenarios, 567–568, 574

review question answers, 851–852
review questions, 641–644
mobile country code (MCC), 576
mobile device management (MDM), 613–626
mobile devices. *See also* laptop and mobile device hardware
characteristics of, 196–199
configuring synchronization, 626–639
malware on, 795–796
troubleshooting
about, 783–784
battery, 784–786
connectivity, 788–792
heat, 784–786
input problems, 786–788
installing, 796–797
malware, 792, 794–796
overheating, 786
performance, 797–799
physical damage, 792–794
power, 784–786
Mobile Equipment Identifier (MEID), 576
mobile hotspots, 567–570
mobile motherboards, 24
mobile network code (MNC), 576
mobile station identifier number (MSIN), 576
modem update, 574
modems, 372
modular power supplies, 126–127
modulation, Wi-Fi and, 492
Molex connector, 184
monitors
about, 144
adjusting display settings, 147–153
exercises, 147–149, 151–153
selecting, 153–155
types of, 144–146
using for input, 146–147
motherboards
about, 5–6, 71
Basic Input/Output System (BIOS), 31–40
bus architecture, 9–10
central processing unit (CPU), 18–24
chipsets, 10–11
complementary metal oxide semiconductor (CMOS), 40–41
components of, 8–41
connector types (expansion slots), 11–15
exam essentials, 71–72
exercises, 34–36
form factors, 6–8
headers, 28–31
for laptops, 212–216
memory slots and cache, 16–18
onboard nonvolatile storage connectors, 25–28
performance-based question answers, 860
performance-based questions, 77
power connectors, 25
processor socket, 18–24
real world scenarios, 14, 21–22
review question answers, 834–835
review questions, 73–76
system board components, 8–41
troubleshooting, 714–716
Unified Extensible Firmware Interface (UEFI), 31–40
mouse, 160–161
mSATA (mini-Serial ATA), 105–106
MS-CHAPv2 (Microsoft Challenge Handshake Authentication Protocol version 2), 592–593
MSIN (mobile station identifier number), 576
multicast address, 420–421
multichannel memory, 52
multicore technology, 46
multifactor authentication (MFA), 458
multifunction devices. *See* printers and multifunction devices
Multi-Link Operation (MLO), 502
multimedia cards, 84–86
multimeter, 736–737
multimode fiber (MMF), 366
multipart paper, 259
multisocket motherboards, 22–23
multitenancy, 671
multithreading CPUs, 47–48
MX (Mail Exchange), 444

N

NaaS (network as a service), 670
naked DSL, 517
Name Server (NS), 444
NAS (network-attached storage), 162, 451, 723
NAT (Network Address Translation), 416, 552
NAT Overload, 418
National Institute of Standards and Technology (NIST), 671
native resolution, 150
near letter quality (NLQ), 260
near-field communication (NFC), 161, 242, 245, 509–510
near-field fingerprint scanner, for laptops and mobile devices, 242
`net` command, 749–753
netbook, 194
`netstat` command, 747–748
Network Address Translation (NAT), 416, 552
network as a service (NaaS), 670
Network Basic Input/Output System (NetBIOS)/NetBIOS over TCP/IP (NetBT), 399–400
network cables, installing, 535–536
network ID, 411
network infrastructure, installing
  about, 530
  connectivity devices, 536–539
  internal NIC, 530–535
  network cables, 535–536
network interface card (NIC), 86–87, 334, 348–353, 530–535
Network layer, in Open Systems Interconnection (OSI) model, 345–346
network services
  about, 439, 480
  authentication, authorization, and accounting (AAA), 457–461
  database servers, 461–465
  Domain Name System (DNS), 442–447
  Dynamic Host Configuration Protocol (DHCP), 448–450
  exam essentials, 480–481
  exercises, 470–472
  file server, 451
  fileshare, 451
  Internet appliances
    about, 465
    load balancing, 468–470
    proxy server, 470–473
    spam gateways, 465–466
    Unified Threat Management (UTM), 466–467
  Internet of Things (IoT) devices
    about, 475
    home security and automation, 477–479
    thermostats, 476
    voice-enabled smart speaker/digital assistants, 479
  legacy/embedded systems, 473–475
  mail server, 452–454
  Network Time Protocol (NTP) server, 465
  performance-based question, 485–486
  performance-based question answers, 865
  print server, 451–452
  real world scenarios, 449
  review question answers, 847–849
  review questions, 482–485
  server roles, 439–442
  syslog, 454–456
  web server, 456–457
network tap, 739, 741
Network Time Protocol (NTP), 465
network-attached storage (NAS), 162, 451, 723
network-based firewalls, 550
networked printers, installing and sharing, 298–309
networking
  about, 325, 381
  bus topology, 341
  client-server resource model, 338–340
  exam essentials, 382–383
  exercises, 371

hardware
about, 348
auxiliary devices, 378–381
cables and connectors, 353–371
connectivity devices, 371–378
network interface card (NIC), 348–353
hybrid topology, 344
IEEE 802 standards, 346–348
local area networks (LANs), 327–328
mesh topology, 343–344
metropolitan area networks (MANs), 330
network resources, 336
network types, 326–327
Open Systems Interconnection (OSI) model, 344–346
peer-to-peer networks, 337–338
performance-based question, 387
performance-based question answers, 863–864
personal area networks (PANs), 329–330
principles of, 325–348
real world scenarios, 335
resource access, 337–340
review question answers, 843–845
review questions, 384–387
ring topology, 342–343
rules of communication, 344–348
servers, 333–334
star topology, 341–342
storage area networks (SANs), 330–332
topologies, 340–344
troubleshooting, 735–763
wide area networks (WANs), 328–329
wireless local area networks (WLANs), 332
workstations, 334–335
NFC (near-field communication), 161, 242, 245, 509–510
NIC (network interface card), 86–87, 334, 348–353, 530–535
NIST (National Institute of Standards and Technology), 671
NLQ (near letter quality), 260
noises, troubleshooting, 705–706
nondedicated servers, 333–334, 440
non-parity memory, 51
Non-Volatile Memory Express (NVMe), 103–104
NS (Name Server), 444
`nslookup` command, 748–749
NTP (Network Time Protocol), 465
number of records, 464

O

OCF (Open Connectivity Foundation), 552
OFDMA (Orthogonal Frequency Division Multiple Access), 494, 501
off-site email application, 678
OLED (organic light-emitting diode), 145–146
OM (optical multimode), 367
onboard nonvolatile storage connectors
about, 25
integrated drive electronics (IDE) drives, 25–26
M.2, 28
Serial ATA (SATA), 26–28
on-demand self-service, 671
ONT (optical network terminal)
about, 372
installing and configuring, 537
Open Connectivity Foundation (OCF), 552
Open Systems Interconnection (OSI) model, 344–346
operating systems
verifying compatibility with, 295–296
vulnerabilities with, 795
optical drives
about, 114–116
for laptops, 225
troubleshooting, 734–735
optical multimode (OM), 367
optical network terminal (ONT)
about, 372
installing and configuring, 537
optical singlemode 1 (OS1), 367
option ROMs, 39

organic light-emitting diode (OLED), 145–146
orientation, printer, 294
Orthogonal Frequency Division Multiple Access (OFDMA), 494, 501
OSI (Open Systems Interconnection) model, 344–346
output, power supply, 121–122
output device. *See* input/output devices
overclocking, 47
overheating, 706–708, 786
ozone, for printers and multifunction devices, 315
ozone filter, for laser printers, 275

P

PaaS (platform as a service), 670
packaging memory, 59–61
packets, 345–346
page printers, 269
page-description language, 289
pairing, 593–605
Palo Alto Research Center (PARC), 160
PANs (personal area networks), 329–330
paper, printer, 310–311
paper feed mechanism, 263–266
paper feed sensors, 266
paper feeder, 265
paper jams, 806, 811–812
paper pickup roller, 273
paper transport assembly, for laser printers, 272–273
paper tray, 265
paper-feeding problems, 806–807
Parallel Advanced Technology Attachment (PATA), 25–26, 95–96, 184
parallel bus, 9–10
parity checking, 51
passive hubs, 374
PAT (Port Address Translation), 418
patch cable, 363
patch panels, 375
`pathping` command, 754–755
pay-as-you-go services, 671
PCB (printed circuit board), 6
PCI (Peripheral Component Interconnect) slots, 11–12
PCI Special Interest Group (PCI-SIG), 14
PCIe (Peripheral Component Interconnect Express)
  about, 10, 101–103
  slots, 12–15
PCL (Printer Control Language), 287–288
peer-to-peer networks, 337–338
Pentium, 21–22
performance
  load balancing and, 469
  of network interface card (NIC), 349
  troubleshooting, 725–730
performance-based question answers
  cables and connectors, 861–862
  cloud computing, 866–867
  cooling systems, 860
  expansion cards, 861
  laptop and mobile device hardware, 862
  memory, 860
  mobile connectivity, 866–867
  motherboards, 860
  network services, 865
  networking, 863–864
  peripherals, 861–862
  power supplies, 861
  printers and multifunction devices, 863
  processors, 860
  small office/home office (SOHO) networks, 866
  storage devices, 861
  Transmission Control Protocol/Internet Protocol (TCP/IP), 864–865
  troubleshooting, 867–868
  wireless networks, 866
performance-based questions
  cables and connectors, 190
  cloud computing, 695
  cooling systems, 77
  expansion cards, 139
  laptop and mobile device hardware, 254
  laptops compared with desktops, 195

memory, 77
mobile application support, 645
mobile connectivity, 645
motherboards, 77
network services, 485–486
networking, 387
peripherals, 190
power supplies, 139
printers and multifunction devices, 321
processors, 77
small office/home office (SOHO) networks, 560
storage devices, 139
Transmission Control Protocol/Internet Protocol (TCP/IP), 435
troubleshooting, 770, 831
virtualization, 695
wireless networks, 560
Peripheral Component Interconnect Express (PCIe)
about, 10, 101–103
slots, 12–15
Peripheral Component Interconnect (PCI) slots, 11–12
peripherals
about, 143, 184
audio devices, 157–158
cable and connectors for, 163–172
exam essentials, 185
exercises, 147–149, 151–153
input and output devices, 159–162
performance-based question, 190
performance-based question answers, 861–862
review question answers, 838–839
review questions, 186–189
storage devices, 162
video devices, 144–157
permissions, 460
personal area networks (PANs), 329–330
personal computer (PC), 5
Personal System/2 (PS/2) connectors, 173
PGA (pin grid array) arrangement, 18–19
physical device damage, troubleshooting, 792–794
Physical layer, in Open Systems Interconnection (OSI) model, 346
physical privacy/security, for laptops, 241–244
physical rate, Wi-Fi and, 492
pickup rollers, 263–264
pickup stepper motor, 264
piconet, 330
piezoelectric effect, 46
pin grid array (PGA) arrangement, 18–19
`ping` command, 745–746
pinwheel, 716
pixel density, 155
pixels, dead, 779–780
plain old telephone system (POTS), 87–88, 517
platform as a service (PaaS), 670
plenum-rate, 354
PoE (Power over Ethernet), 379–381, 513
point stick, for laptops, 229–230
pointing devices, for laptops, 226–231
point-to-multipoint, 521
polyvinyl chloride (PVC), 354
Port 20/21-File Transfer Protocol (FTP), 396–397
Port 22-Secure Shell (SSH), 397
Port 23-Telnet, 397–398
Port 25-Simple Mail Transfer Protocol (SMTP), 398
Port 53-Domain Name System (DNS), 398–399
Port 67/68-Dynamic Host Configuration Protocol (DHCP), 399
Port 80-Hypertext Transfer Protocol (HTTP), 399
Port 110-POP3, 399
Port 137-139-Network Basic Input/Output System (NetBIOS)/NetBIOS over TCP/IP (NetBT), 399–400
Port 143-Internet Message Access Protocol (IMAP), 400
Port 161/162-Simple Network Management Protocol (SNMP), 400–401
Port 389-Lightweight Directory Access Protocol (LDAP), 401

Port 443-Hypertext Transfer Protocol Secure (HTTPS), 401–402
Port 445-Server Message Block (SMB)/ Common Internet File System (CIFS), 402
Port 3389-Remote Desktop Protocol (RDP), 402–403
Port Address Translation (PAT), 418
port flapping, 760, 761
port mirroring, 376
port replicators, 247–249
port triggering, 551
portability, laptops compared with desktops, 195
ports
about, 162, 267
serial, 172–173
setting up, 244–249
Ports tab (Printer Properties), 292
Post Office Protocol 3 (POP3), 399, 454
PostScript (PS), 287–288
POTS (plain old telephone system), 87–88, 517
power
troubleshooting, 807, 809
troubleshooting issues, 784–786
power adapters, for laptops, 233, 235–236
power button
on motherboard headers, 30
Universal Serial Bus (USB), 165–166
power circuits, 267
power connectors, 25
power light, on motherboard headers, 30
Power over eSATA, 183
Power over Ethernet (PoE), 379–381, 513
power supply
about, 81, 119–120, 132
AC adapters as, 131–132
connectors, 123–126
energy efficiency, 122–123
exam essentials, 132–133
exercises, 131
input, 120–121
modular, 126–127
output, 121–122
performance-based question answers, 861
performance-based questions, 139
real world scenario, 122
redundant, 127–130
replacing, 130–131
review question answers, 836–837
review questions, 134–138
troubleshooting, 719–722
wattage ratings, 121–122
power supply exhaust fan, 63
power-on self-test (POST) routine
about, 39, 59
troubleshooting, 710, 712–713
preferred roaming list (PRL) updates, 576
Presentation layer, in Open Systems Interconnection (OSI) model, 345
PRI (product release instruction) updates, 576
print jobs, managing, 817–819
print options, managing, 821–823
print quality, troubleshooting, 800–801, 803–805
print queue, 290
print server, 333, 451–452
print spooler, managing, 820–821
printed circuit board (PCB), 6
printer banks, 512
printer control circuits, 267
Printer Control Language (PCL), 287–288
printer controller assembly, for laser printers, 275
printer driver, 816–817
printer jams, 802
printer paper, 310–311
printer ribbon, 258, 311
printers and multifunction devices
about, 257, 315–316
calibrating, 288
daisy-wheel, 258–259
dot-matrix, 259–260
exam essentials, 316–317
exercises, 267, 297, 299–300, 301–302, 308–309, 314
imaging processes, 257–282
impact, 258–260
inkjet, 260–269

installing, 282–315
laser, 269–281
local network, 299–303
maintaining, 282–315
network scan services, 306–309
network security, 303–306
networked, 298–309
performance-based question, 321
performance-based question answers, 863
print technologies, 257–282
real world scenarios, 289, 297–298, 312
review question answers, 841–843
review questions, 318–321
thermal, 281–282
troubleshooting
about, 799–800
impact, 800–803
inkjet, 803–807
laser, 807–817
managing print jobs, 817–819
managing print options, 821–823
managing print spooler, 820–821
printhead, 258, 261–262
printhead alignment, 804–805
printhead carriage, 262–263
privacy screens, for laptops, 243–244
private cloud, 669
private IP addresses, 416–418
private side, 549
PRL (preferred roaming list) updates, 576
Process/Application layer protocols, 396–403
processing step, in electrophotographic (EP) imaging process, 276, 808
processor socket, 18–21
processor throttling, 216
processors
about, 5–6, 42–43, 71
CPU architecture, 43–46
CPU characteristics, 46–49
exam essentials, 71–72
for laptops, 216–217
performance-based question answers, 860
performance-based questions, 77
real world scenarios, 48–49
review question answers, 834–835
review questions, 73–76
product release instruction (PRI) updates, 576
programmable ROM (PROM), 59
projection systems
about, 155–157
burnt-out bulbs, 776
troubleshooting, 782
propagation delay, 521
proxy server, 470–473
PS (PostScript), 287–288
PS/2 (Personal System/2) connectors, 173
public cloud, 668–669
public IP addresses, 416–418
public side, 549
punchdown blocks, 362–363
punchdown tool, 739
push protocol, 398
PVC (polyvinyl chloride), 354

Q

quad A (AAAA), 444
quad-channel memory, 52
Quadrature Amplitude Modification (QAM), 494
Quadrature Phase Shift Keying (QPSK), 494
quality of construction, laptops compared with desktops, 195
quality of service (QoS), 376
QWERTY keyboard, 159–160

R

radio firmware
about, 574
update, 574
radio frequency identification (RFID), 508–509
radio frequency networking standards, 508–510
RADIUS (Remote Authentication Dial-In User Service), 458

RAID 1+0, 110, 111
RAID (Redundant Array of Independent Disks) arrays
  about, 109–111
  troubleshooting, 722–735
random access memory (RAM) slots, 16
random shutdowns, 708
ranges, Wi-Fi, 498–499
rapid elasticity, 671
RARP (Reverse ARP), 394
Raster Image Processor (RIP), 276
raster line, 276
rasterizing, 275
RDP (Remote Desktop Protocol), 402–403
read-only memory (ROM), 54, 58–59
real world scenarios
  laptop and mobile device hardware, 207–208, 235
  mobile connectivity, 567–568, 574
  motherboards, 14, 21–22
  network services, 449
  networking, 335
  power supplies, 122
  printers and multifunction devices, 289, 297–298, 312
  processors, 48–49
  small office/home office (SOHO) networks, 522–523, 535, 546–547, 553
  Transmission Control Protocol/Internet Protocol (TCP/IP), 401–402, 417, 419
  troubleshooting, 703, 707–708, 720, 758–759, 790, 798–799, 812
  virtualization, 660
real-time operating systems (RTOSs), 574
rear exhaust fan, 63
reassembling laptops, 199–202
receivers, GPS, 608
Reduced Instruction Set Computer (RISC), 43–45
redundancy, 376
Redundant Array of Independent Disks (RAID) arrays
  about, 109–111
  troubleshooting, 722–735
redundant power supply (RPS), 127–130
refresh rate, 147–149, 154
refresh signal, 54
registered jack (RJ) connector, 361–362
registered ports, 395
registration roller, 273
reliability, load balancing and, 469
Remote Authentication Dial-In User Service (RADIUS), 458
Remote Desktop Protocol (RDP), 402–403
removable storage/media
  about, 111
  flash memory, 111–113
  hot-swappable devices, 113–114
  optical drives, 114–116
removing storage devices, 116–118
repeaters, 373
replacing
  consumables, 310–313
  power supplies, 130–131
reservations, 448–449
reset button, on motherboard headers, 30
resistive touchscreens, 146
resolution, 149–150, 154
resource pooling, 671
resources
  networking, 336–337
  requirements for virtualization, 656–659
Reverse ARP (RARP), 394
review question answers
  cables and connectors, 838–839
  cloud computing, 853–854
  cooling systems, 834–835
  expansion cards, 836–837
  laptop and mobile device hardware, 839–841
  memory, 834–835
  mobile application support, 851–852
  mobile connectivity, 851–852
  motherboards, 834–835
  network services, 847–849
  networking, 843–845
  peripherals, 838–839
  power supplies, 836–837

printers and multifunction devices, 841–843
processors, 834–835
small office/home office (SOHO) networks, 849–850
storage devices, 836–837
Transmission Control Protocol/Internet Protocol (TCP/IP), 845–847
troubleshooting, 854–858
virtualization, 853–854
wireless networks, 849–850
review questions
cables and connectors, 186–189
cloud computing, 691–694
cooling systems, 73–76
expansion cards, 134–138
laptop and mobile device hardware, 250–253
memory, 73–76
mobile application support, 641–644
mobile connectivity, 641–644
motherboards, 73–76
network services, 482–485
networking, 384–387
peripherals, 186–189
power supplies, 134–138
printers and multifunction devices, 318–321
processors, 73–76
small office/home office (SOHO) networks, 556–559
storage devices, 134–138
Transmission Control Protocol/Internet Protocol (TCP/IP), 432–435
troubleshooting, 766–769, 827–830
virtualization, 691–694
wireless networks, 556–559
RFID (radio frequency identification), 508–509
ribbon cartridge, 311
right-click, 161
ring topology, 342–343
RIP (Raster Image Processor), 276
RISC (Reduced Instruction Set Computer), 43–45
riser cards, 15, 212–213
RJ (registered jack) connector, 361–362
ROM (read-only memory), 54, 58–59
ROM BIOS chip. *See* Basic Input/Output System (BIOS)
routers, 377–378
Routing and Remote Access Service (RRAS), 458
RPS (redundant power supply), 127–130
RTOSs (real-time operating systems), 574

S

SaaS (software as a service), 670
SAE (Simultaneous Authentication of Equals), 545
SAM (Slot Availability Masking), 506–507
sandbox, 651
SANs (storage area networks), 330–332, 451
SAS (Serial Attached SCSI), 28, 104–105
SASI (Shugart Associates System Interface) standard, 28
SATA (Serial Advanced Technology Attachment)
about, 10, 26–28, 100–101
connectors, 182–183
power connectors, 125–126
satellite Internet, 521–523
SC (subscriber connector), 370
SCA (Single Connector Attachment), 114
SCADA (supervisory control and data acquisition), 474–475
scalability
of database servers, 463
load balancing and, 469
scan line, 276
scan services, for networked printers, 306–309
scan to cloud, 308
scan to email, 307
scan to folder, 308
scatternet, 330
scheduled maintenance, 313
scopes, 448–449

screen, for laptops, 208–209
screen resolution, 149–150, 154
screened subnet, 379, 441, 549
SCSI (Small Computer System Interface), 10, 28, 104–105
SDR SDRAM (single data rate SDRAM), 54, 55
SDRAM (synchronous DRAM), 54
SDSL (symmetric DSL), 514–515
sealing, 36
secondary storage, 198–199
Secure Boot, 37–39
Secure Digital (SD) format, 113
secure element identifier (SEID), 577
Secure Shell (SSH), 397
Secure Sockets Layer (SSL), 401, 625
secured prints, 304
security
  BIOS/UEFI and, 36–39
  of networked printers, 303–306
  requirements for virtualization, 659–667
SEID (secure element identifier), 577
selecting
  internal network connections, 526–530
  Internet connections, 514–526
  monitors, 153–155
Self-Monitoring, Analysis, and Reporting Technology (S.M.A.R.T.), 730–733
Sender Policy Framework (SPF), 445
sending data, network interface card (NIC) and, 349–350
separation pads, 264, 806
Serial Advanced Technology Attachment (Serial ATA/SATA)
  about, 10, 26–28, 100–101
  connectors, 182–183
  power connectors, 125–126
Serial Attached SCSI (SAS), 28, 104–105
serial bus, 10
serial ports, 172–173
Server Message Block (SMB), 308, 402
server roles, 439–442
server-based model, 338–340
servers, 333–334
service-set identifier (SSID), 490, 757
servicing mobile devices, 196–197
Session layer, in Open Systems Interconnection (OSI) model, 345
settings, adjusting for displays, 147–153
setup
  accessories, 244–249
  cloud servers, 678–688
  ports, 244–249
severity level, 455
SFA (single-factor authentication), 458
SFF (small form factor), 8
shared resources, 671
ShareNet, 327–328
sharing networked printers, 298–309
Sharing tab (Printer Properties), 292
shielded twisted pair (STP), 358–359
shingled magnetic recording (SMR), 98–99
Shugart Associates System Interface (SASI) standard, 28
SIM (subscriber identity module), 575
Simple Mail Transfer Protocol (SMTP), 393, 398, 454
Simple Network Management Protocol (SNMP), 400–401
Simultaneous Authentication of Equals (SAE), 545
simultaneous multithreading (SMT), 47
Single Connector Attachment (SCA), 114
single data rate SDRAM (SDR SDRAM), 54, 55
single point of failure, 341
single-channel memory, 52
single-core technology, 46
single-factor authentication (SFA), 458
single-mode fiber (SMF), 366
sizing issues, 777–778
Slot Availability Masking (SAM), 506–507
Small Computer System Interface (SCSI), 10, 28, 104–105
small form factor (SFF), 8
small office/home office (SOHO) networks
  about, 489, 511, 554
  cable, 518–519
  cellular
    about, 523

fifth generation (5G), 524–525
fourth generation (4G), 523–524
choosing internal network connections
about, 526–527
wired network connections, 527–529
wireless network connections, 529–530
choosing Internet connections, 514
configuring wireless routers and access points
about, 539–547
wireless channels, 547–548
wireless router services, 548–553
digital subscriber line (DSL), 514–517
exam essentials, 554–555
exercises, 529–531
fiber-optic Internet, 519–520
installing network infrastructure
about, 530
installing and configuring connectivity devices, 536–539
installing internal NIC, 530–535
installing network cables, 535–536
performance-based question answers, 866
performance-based questions, 560
planning keys, 511–513
real world scenarios, 522–523, 535, 546–547, 553
review question answers, 849–850
review questions, 556–559
satellite, 521–523
wireless Internet service provider (WISP), 525–526
small outline dual inline memory module (SODIMM), 16, 61, 217–220
S.M.A.R.T. (Self-Monitoring, Analysis, and Reporting Technology), 730–733
SMB (Server Message Block), 308, 402
SMF (single-mode fiber), 366
smoke, troubleshooting, 708–709
SMP (symmetric multiprocessing), 47
SMR (shingled magnetic recording), 98–99
SMT (simultaneous multithreading), 47
SMTP (Simple Mail Transfer Protocol), 393, 398, 454
smudging, troubleshooting in printing, 815–816
sneakernet, 326
SNMP (Simple Network Management Protocol), 400–401
SOA (Start of Authority), 444
SODIMM (small outline dual inline memory module), 16, 61, 217–220
software as a service (SaaS), 670
software commands, 741–755
SOHO (small office/home office) networks
about, 489, 511, 554
cable, 518–519
cellular
about, 523
fifth generation (5G), 524–525
fourth generation (4G), 523–524
choosing internal network connections
about, 526–527
wired network connections, 527–529
wireless network connections, 529–530
choosing Internet connections, 514
configuring wireless routers and access points
about, 539–547
wireless channels, 547–548
wireless router services, 548–553
digital subscriber line (DSL), 514–517
exam essentials, 554–555
exercises, 529–531
fiber-optic Internet, 519–520
installing network infrastructure
about, 530
installing and configuring connectivity devices, 536–539
installing internal NIC, 530–535
installing network cables, 535–536
performance-based question answers, 866
performance-based questions, 560
planning keys, 511–513
real world scenarios, 522–523, 535, 546–547, 553
review question answers, 849–850
review questions, 556–559
satellite, 521–523
wireless Internet service provider (WISP), 525–526
solenoid, 258

solid-state drives (SSDs)
about, 99–100
communication interfaces, 100–105
form factors, 105–108
for laptops, 220–221, 224
sound cards, 84–85
sounds, troubleshooting, 723–724
spam gateways, 465–466
spam management, 444–445
Spanning Tree Protocol (STP), 423
spatial streams, Wi-Fi and, 492
speakers, 247
speed
of CPUs, 46–47
of database servers, 463
troubleshooting, 717–719
SPF (Sender Policy Framework), 445
spindle speeds, for HDD systems, 96–98
splitter, 356–357, 358
square connector, 370
SRAM (static random access memory), 54, 58
sRGB, 155
SSDs (solid-state drives)
about, 99–100
communication interfaces, 100–105
form factors, 105–108
for laptops, 220–221, 224
SSH (Secure Shell), 397
SSID (service-set identifier), 490, 757
SSL (Secure Sockets Layer), 401, 625
ST (straight tip), 370
stabilizer bar, 263
standards
for fiber-optic cables, 368–369
Thunderbolt, 170–171
Universal Serial Bus (USB), 164–166
star topology, 341–342
Start of Authority (SOA), 444
STARTTLS, 625
static IP address, 403, 449
static random access memory (SRAM), 54, 58
static-charge eliminator strip, 274
stepper motor
about, 263
troubleshooting, 802–803, 807
storage, for laptops, 220–225
storage area networks (SANs), 330–332, 451
storage devices
about, 81, 94, 132
exam essentials, 132–133
exercises, 117–119
hard disk drive (HDD) systems, 94–99
installing, 116, 118–119
performance-based question answers, 861
performance-based questions, 139
as peripherals, 162
Redundant Array of Independent Disks (RAID), 109–111
removable, 111–116
removing, 116–118
review question answers, 836–837
review questions, 134–138
solid-state drives (SSDs), 99–108
STP (shielded twisted pair), 358–359
STP (Spanning Tree Protocol), 423
straight tip (ST), 370
stripe set with parity (RAID 5), 109–110, 111
stylus, 245–246
subnet masks, 411, 417
subscriber connector (SC), 370
subscriber identity module (SIM), 575
super scope, 425
SuperDuper, 224
supervisory control and data acquisition (SCADA), 474–475
suspicious links, 795
switches
about, 376–377
installing and configuring, 537
symmetric DSL (SDSL), 514–515
symmetric multiprocessing (SMP), 47
synchronization, configuring for mobile devices, 626–639
synchronization app, 638
synchronous DRAM (SDRAM), 54
syslog, 454–456
system board. *See* motherboards
system clock, 46

T

TCP printing, 299–301
TCP/IP (Transmission Control Protocol/Internet Protocol)

about, 391–392, 423, 429–430
exam essentials, 430–431
exercises, 403–409, 427–429
IP addressing
about, 403–409
Automatic Private IP Addressing (APIPA), 418–420
binary values, 410
class inter-domain routing (CIDR), 414–416
components of, 411–412
IPv4 address classes, 412–414
public *vs.* private, 416–418
IPv6
about, 420–421
working with, 421–422
performance-based question, 435
performance-based question answers, 864–865
real world scenarios, 401–402, 417, 419
review question answers, 845–847
review questions, 432–435
structure
about, 392–393
Host-to-Host layer protocols, 394–396
Internet layer protocols, 394
process/application layer protocols, 396–403
virtual local area networks (VLANs)
about, 423–425
virtual private networks (VPNs), 425–429
Telnet, 397–398
temperature monitoring, 34
Temporal Key Integrity Protocol (TKIP), 544
Terminal Access Controller Access-Control System Plus (TACACS+), 458
test development, 652
test pages, printing, 290, 295, 821
tethering, 245, 567–570
Text record (TXT), 444
thermal pad, 66–67
thermal paper, 281
thermal paste, 66
thermal printers, 281–282
thermal transfer printers, 282
thermal wax-transfer printers, 282
thermostats, 476
third generation (3G), 565
three-pronged firewall, 441
Thunderbolt connectors, 170–172
TKIP (Temporal Key Integrity Protocol), 544
TLS (Transport Layer Security), 401, 625
TN (twisted nematic), 145
toggling the display, 775
toner, 269
toner cartridge, 270–271, 311–313, 812–813
toner probe, 738–739
tools
for disassembling and reassembling laptops, 200–201
for network troubleshooting, 735–755
topologies, network, 340–344
touch pen, 245–246
TouchFLO, 198
touchpads, 227–228
touchscreen, 146, 230–231
TPM (Trusted Platform Module), 36–37
`tracert` command, 753–754
track point, 229–230, 247
trackballs, 161, 226
trackpad, 227–228, 247
tractor feed, 258
transfer corona assembly, 273–274, 813
transfer rates, 116
transferring step, in electrophotographic (EP) imaging process, 278–279, 808
transformer, 267, 807
Transmission Control Protocol/Internet Protocol (TCP/IP)
about, 391–392, 423, 429–430
exam essentials, 430–431
exercises, 403–409, 427–429
IP addressing
about, 403–409
Automatic Private IP Addressing (APIPA), 418–420
binary values, 410
class inter-domain routing (CIDR), 414–416

components of, 411–412
IPv4 address classes, 412–414
public *vs.* private, 416–418
IPv6
about, 420–421
working with, 421–422
performance-based question, 435
performance-based question answers, 864–865
real world scenarios, 401–402, 417, 419
review question answers, 845–847
review questions, 432–435
structure
about, 392–393
Host-to-Host layer protocols, 394–396
Internet layer protocols, 394
process/application layer protocols, 396–403
virtual local area networks (VLANs)
about, 423–425
virtual private networks (VPNs), 425–429
Transport layer, in Open Systems Interconnection (OSI) model, 345
Transport Layer Security (TLS), 401, 625
triangulation, 607
triple-channel memory, 52
troubleshooting
about, 700, 704–705, 763, 773–774, 823–824
best practices
about, 701
structured methodology, 702–704
tips and tricks, 701–702
BIOS/UEFI problems, 710–712
blank screen, 715
burning smells and smoke, 708–710
cable problems, 714–715
capacitor swelling, 715–716
CPU problems, 714–716
display issues, 781–783
drives, 722–735
exam essentials, 763–765, 824–826
exercises, 721–722, 726–729, 743–745, 751–753, 820–821
hardware
about, 705
grinding and other noises, 705–706
overheating, 706–708
intermittent device failure, 710
I/O port problems, 714–715
memory problems, 716–719
mobile device issues
about, 783–784
battery, 784–786
connectivity, 788–792
heat, 784–786
input problems, 786–788
installing, 796–797
malware, 792, 794–796
overheating, 786
performance, 797–799
physical damage, 792–794
power, 784–786
motherboard problems, 714–716
networking problems, 735–763
performance-based question, 770, 831
performance-based question answers, 867–868
POST routine problems, 712–713
power supply problems, 719–722
printers
about, 799–800
impact, 800–803
inkjet, 803–807
laser, 807–817
managing print jobs, 817–819
managing print options, 821–823
managing print spooler, 820–821
RAID arrays, 722–735
real world scenarios, 703, 707–708, 720, 758–759, 790, 798–799, 812
review question answers, 854–858
review questions, 766–769, 827–830
video image problems, 776–781
video input problems, 775–776
TRS connectors, 157–158
Trusted Platform Module (TPM), 36–37
20-pin to 24-pin motherboard adapter, 124
twisted nematic (TN), 145

twisted pair cable, for networking, 358–365
2-in-1 devices, 231
TXT (Text record), 444
Type 1 hypervisor, 652–653
Type 2 hypervisor, 655

U

ubiquitous access, 671
UDP (User Datagram Protocol), 393, 394
Ultra-Reliable Low-Latency Communications (URLLC), 566
unicast address, 420–421
unidirectional, 14
Unified Extensible Firmware Interface (UEFI)
  about, 32–33
  configuration and monitoring, 33–34
  exercises, 34–36
  flashing the system, 39–40
  power-on self-test (POST) routine, 39
  security and encryption, 36–39
  troubleshooting, 710–712
Unified Threat Management (UTM), 466–467
uninterruptible power supply (UPS), 128–130
Universal Plug and Play (UPnP), 552–553
Universal Serial Bus (USB)
  about, 10
  cables and connectors, 164–169
  expansion cards, 88–89
  flash drives, 112
  hub, 164
  for laptops and mobile devices, 244
  permissions, 34
  ports on motherboard headers, 31
  for printers and multifunction devices, 284
universal voltage, 121
unmanaged switch, 376
unsecure networks, 796
unshielded twisted pair (UTP), 358–359
updating
  device drivers, 89–92
  driver software, 286–288
  firmware, 288
UPnP (Universal Plug and Play), 552–553
UPS (uninterruptible power supply), 128–130
URLLC (Ultra-Reliable Low-Latency Communications), 566
USB (Universal Serial Bus)
  about, 10
  cables and connectors, 164–169
  expansion cards, 88–89
  flash drives, 112
  hub, 164
  for laptops and mobile devices, 244
  permissions, 34
  ports on motherboard headers, 31
  for printers and multifunction devices, 284
USB Battery Charging, 165
USB Power Delivery (PD), 165–166
USB Type-C (USB-C), 168
user authentication, 303
User Datagram Protocol (UDP), 393, 394
UTM (Unified Threat Management), 466–467
UTP (unshielded twisted pair), 358–359

V

vampire tap, 356
VDC (volts DC), 120
VDI (virtual desktop infrastructure), 658
verifying compatibility, 295–296
versions, Bluetooth, 505–507
vertical alignment (VA), 145
very high bit-rate DSL (VDSL), 514–515
VESA (Video Electronics Standards Association), 180
VGA (Video Graphics Array) connector, 174–175
video
  cables and connectors
    about, 173–174
    Digital Visual Interface (DVI), 175–177
    DisplayPort, 180–181
    High-Definition Multimedia Interface (HDMI), 177–180
    Video Graphics Array (VGA), 174–175

devices
about, 144
monitors, 144–155
projection systems, 155–157
troubleshooting
image issues, 776–781
input issues, 775–776
video capture cards, 85–86
video cards
about, 82–83, 144
cooling, 70
for laptops, 207–208
Video Electronics Standards Association (VESA), 180
Video Graphics Array (VGA) connector, 174–175
video random-access memory (VRAM), 82
virtual application, 675
virtual desktop infrastructure (VDI), 658
virtual local area networks (VLANs), 376, 423–425
virtual machine manager (VMM), 652
virtual machines (VMs), 650–652
virtual networks
about, 423
virtual local area networks (VLANs), 423–425
virtual private networks (VPNs), 425–429
virtual NIC, 658
virtual private networks (VPNs), 425–429
virtual sandbox, 651
VirtualBox, 661–667
virtualization
about, 649–650, 688–689
client-side requirements
about, 655, 656
emulator requirements, 659
resource requirements, 656–659
security requirements, 659–667
containers, 655–656
exam essentials, 689–690
exercises, 653–654, 661–667
hypervisor, 652–655
performance-based question, 695
purpose of virtual machines (VMs), 650–652
real world scenarios, 660
review question answers, 853–854
review questions, 691–694
Virtualization Technology (VT)
about, 657
for CPUs, 49
VLANs (virtual local area networks), 376, 423–425
VMs (virtual machines), 650–652
VNN (virtual machine manager), 652
Voice over Internet Protocol (VoIP), 761
voice-enabled smart speakers/digital assistants, 479
VoIP (Voice over Internet Protocol), 761
volts DC (VDC), 120
volume set, 109
VPNs (virtual private networks), 425–429
VRAM (video random-access memory), 82
VT (Virtualization Technology)
about, 657
for CPUs, 49

W

WANs (wide area networks), 328–329
WAP (wireless access point)
about, 332, 372–373, 504
installing and configuring, 537–539
warm-swappable devices, 114
wattage ratings, 121–122
web server, 456–457
webcams, 208, 247
websites
Andy, 659
Iometer, 729
Laplink PCmover, 224
Macrium Reflect X, 224
Open Connectivity Foundation (OCF), 552
PCIe, 14
SuperDuper, 224
Wi-Fi Alliance, 539

weight, of printer paper, 310
well-known ports, 395
white lines, troubleshooting in printing, 815
wide area networks (WANs), 328–329
Wi-Fi. *See* wireless networks
Wi-Fi Alliance, 539
Wi-Fi Analyzer, 737–738
Wi-Fi antenna, on laptops, 209
Wi-Fi Protected Access (WPA), 544–545
Wi-Fi Protected Access 2 (WPA2), 545
Wi-Fi Protected Access 3 (WPA3), 545
WiMAX, 524
wired network connections, 527–529
wireless access point (WAP)
  about, 332, 372–373, 504
  installing and configuring, 537–539
wireless channels, 547–548
Wireless Internet service provider (WISP), 525–526
wireless local area networks (WLANs), 332
wireless networks
  about, 489–490, 554
  Bluetooth
    about, 505
    devices, 507–508
    versions, 505–507
  connections, 529–530
  802.11 networking
    about, 490–491
    devices, 504–505
    gigabit standards, 499–504
    legacy standards, 492–499
    terminology for, 491–492
  establishing Wi-Fi connectivity, 577–593
  exam essentials, 554–555
  performance-based question answers, 866
  performance-based questions, 560
  for printers and multifunction devices, 284–285
  radio frequency standards
    about, 508
    near-field communication (NFC), 509–510
    radio frequency identification (RFID), 508–509
  review question answers, 849–850
  review questions, 556–559
wireless personal area network (WPAN), 329–330
wireless router services, 548–553
wireless routers
  configuring, 539–553
  installing and configuring, 537–539
wiring standards, for twisted pair cables, 363–364
WISP (Wireless Internet service provider), 525–526
WLANs (wireless local area networks), 332
workgroups, 337–338
workstations, 334–335
WPA (Wi-Fi Protected Access), 544–545
WPA 2 (Wi-Fi Protected Access 2), 545
WPA 3 (Wi-Fi Protected Access 3), 545
WPAN (wireless personal area network), 329–330

X

x64/x86, 44
XaaS (Anything/everything as a service), 670

Z

zero configuration networking, 419
zero insertion force (ZIF) sockets, 20
zone file, 443

Online Test Bank

To help you study for your CompTIA A+ certification exams, register to gain one year of FREE access after activation to the online interactive test bank—included with your purchase of this book!

To access our learning environment, simply visit `www.wiley.com/go/sybextestprep`, follow the instructions to register your book, and instantly gain one year of FREE access after activation to:

- All of the chapter questions from this book, so you can practice in a timed and graded setting
- A set of nearly 400 flashcards
- A searchable glossary

Do you need more practice? Check out these resources:

CompTIA A+ Complete Practice Tests: Core 1 Exam 220-1201 and Core 2 Exam 220-1202, Fourth Edition by Audrey O'Shea (ISBN 9781394330331) With hundreds of practice questions for each domain and with additional complete practice exams, the *Practice Tests* series a great way to build your confidence and readiness for exam day.

CompTIA A+ Complete Review Guide: Core 1 Exam 220-1201 and Core 2 Exam 220-1202, Sixth Edition by Troy McMillan (ISBN 9781394330782) The *Review Guide* series is a must-have companion resource to the *CompTIA A+ Complete Study Guide*. It includes practical examples and insights drawn from the real-world experience of PC and hardware technicians, as well as A+ certification exam highlights and concise end-of-chapter reviews.

SYBEX
A Wiley Brand

Your COMPLETE resource for the UPDATED CompTIA A+ Core 1 and Core 2 exams

Building on the popular Sybex self-study approach, this Fifth Edition of bestselling *CompTIA® A+® Complete Study Guide: Core 1 Exam 220-1101 and Core 2 Exam 220-1102*, is revised and updated to cover major advances in technology, including mobile, cloud, networking, and security. The Study Guide includes practical examples and insights drawn from real-world experience, as well as exam highlights and end-of-chapter review. Readers also have access to additional learning resources, including the Sybex interactive online learning environment and test bank with hundreds of review questions, practice exams, electronic flashcards, and a searchable PDF of a Glossary of the most important terms readers will need to understand.

The book provides 100% coverage of the CompTIA A+ Core 1 and Core 2 exam objectives, and verifies the successful candidate has the knowledge and job skills needed for:

- Installing, configuring, and maintaining computer equipment, mobile devices, and software for end users
- Servicing components based on customer requirements
- Understanding networking basics and applying basic cybersecurity methods to mitigate threats
- Properly and safely diagnosing, resolving, and documenting common hardware and software issues
- Applying troubleshooting skills and providing customer support using appropriate communication skills
- Understanding the basics of scripting, cloud technologies, virtualization, and multi-OS deployments in corporate environments

ABOUT CompTIA A+ CERTIFICATION

The CompTIA A+ certification is the industry standard for computer technicians. Exams measure a candidate's competency with various hardware components and software operating systems. A+ certification is the most popular entry-level IT certification, especially for those switching to IT from another career.

ABOUT THE AUTHORS

QUENTIN DOCTER, A+, Network+, IT Fundamentals+, Cloud Essentials+, is an IT consultant who started in the industry in 1994 working in tech support for Packard Bell. Since then, he's worked as a tech and network support specialist, trainer, consultant, and webmaster. During his career, he has achieved certifications from CompTIA, Microsoft, Cisco, Novell, and Sun Microsystems. He has written over a dozen books for Sybex, including the *CompTIA IT Fundamentals+ Study Guide* and the *CompTIA Cloud Essentials+ Study Guide*.

JON BUHAGIAR, A+, Network+, is an information technology professional with two decades of experience in higher education. During the past 22 years he has been responsible for Network Operations at Pittsburgh Technical College and lead several projects, including virtualization (server and desktop), VoIP, Microsoft 365, and many other projects supporting the quality of education at the college. He has achieved several certifications from CompTIA, Cisco, and Microsoft, and taught many of the certification paths. He is the author of several books, including Sybex's *CompTIA Network+ Review Guide*.

www.sybex.com
www.wiley.com/go/sybextestprep

Cover Design: Wiley
Cover Image: © Getty Images Inc./Jeremy Woodhouse

SYBEX
A Wiley Brand

Also available as an e-book

Interactive learning environment

Take your exam prep to the next level with Sybex's superior interactive online study tools. To access our learning environment, simply visit **www.wiley.com/go/sybextestprep**, register your book to receive your unique PIN, and instantly gain one year of FREE access after activation to:

- **Interactive test bank with 8 practice exams** to help you identify areas where further review is needed.
- **Over 500 electronic flashcards** to reinforce learning and last-minute prep before the exam
- **Comprehensive glossary** in PDF format gives you instant access to the key terms so you are fully prepared
- **Over 2 hours of recorded audio instruction** from the authors to help you study for the certification exam with confidence

CATEGORY
COMPUTERS/Certification Guides

ISBN 978-1-119-86291-8
56000
9 781119 862918